Analyse van biologische beelden

Biological Image Analysis

Jonas De Vylder

Promotor: prof. dr. ir. W. Philips
Proefschrift ingediend tot het behalen van de graad van
Doctor in de Ingenieurswetenschappen: Computerwetenschappen

Vakgroep Telecommunicatie en Informatieverwerking
Voorzitter: prof. dr. ir. H. Bruneel
Faculteit Ingenieurswetenschappen en Architectuur
Academiejaar 2014 - 2015

UNIVERSITEIT
GENT

ISBN 978-90-8578-737-2
NUR 959, 954
Wettelijk depot: D/2014/10.500/83

Members of the jury

prof. dr. Daniel Ochoa Donoso (Escuela Superior Politécnica del Litoral, ESPOL University)
prof. dr. Dirk Inzé (Vlaams instituut voor biotechnologie, Ghent University)
prof. dr. Nele Noels (Ghent University, secretary)
prof. dr. Wilfried Philips (Ghent University, supervisor)
prof. dr. Jan Van Campenhout (Ghent University, chairman)
prof. dr. Dominique Van Der Straeten (Ghent University)
dr. Danilo Babin (Ghent University)
dr. Christopher Guérin (Vlaams instituut voor biotechnologie, Ghent University)
dr. Filip Rooms (Vigo)

Affiliations

Research Group for Image Processing and Interpretation (IPI)
iMinds
Department of Telecommunications and Information Processing (TELIN)
Faculty of Engineering and Architecture
Ghent University

Sint-Pietersnieuwstraat 41
B-9000 Ghent
Belgium

Acknowledgement

This work would not have been possible without the help and support of many people.

First and foremost, I would like to thank my family and friends. Without your continuous support and encouragement, this thesis would never have been written. Kristien and Fabrice, you both even experienced the horror of proof reading some chapters, for which I am very grateful.

I want to express my gratitude to my supervisor professor Wilfried Philips who not only gave me the opportunity to do this research, but also gave me the freedom to investigate and develop my own research topics and of course who also gave me plenty of feedback and discussions. Furthermore, I would like to thank the members of my Ph.D. jury, prof. Jan Van Campenhout, prof. Dominique Van Der Straeten, prof. Dirk Inzé, dr. Christopher Guérin, dr. Nele Noels, dr. Danilo Babin, dr. Filip Rooms, and dr. Daniel Ochoa Donoso for being in the jury, for thoroughly reading this dissertation, for the interesting discussions and for their useful feedback.

Biological image analysis is intrinsically an interdisciplinary research topic, fortunately I could rely on the great expertise of many other researchers. My special thanks goes to Winnok De Vos (Cell Systems & Imaging, Ghent University); Bruce McNaughton and his team (Polaris Brain Dynamics research group, The University of Lethbridge, Canada); Cynthia Lambrecht, Mado Vandewoestyne, Trees Lepez and Dieter Deforce (Laboratory for Pharmaceutical Biotechnology, Ghent University); Dominique Van Der Straeten, Filip Vandenbussche, Yuming Hu, Dajo Smet, Pieter Callebert and Laury Chaerle (Laboratory of Functional Plant Biology, Ghent University); Stijn Dhondt and Dirk Inzé (Department of Plant Systems Biology, Vlaams instituut voor biotechnologie, Ghent University); Saskia Lippens, Yvan Saeys, Anneke Kremer, Michiel Krols, Bob Asselbergh, Sophie Janssens and Chris Guérin (Inflammation Research Center and Bio-imaging Core, Vlaams instituut voor biotechnologie, Ghent University). It is often said that the key to successfully finishing a Ph.D. is to keep yourself motivated. It were the challenges and numerous interesting applications that you all provided that kept me motivated. I always could rely on your expertise, your willingness to explain in layman

terms and your enthusiasm to design and do new experiments. I've sincerely enjoyed our fruitful collaboration.

This work is not only the result of collaboration between other research groups, but also due to the wonderful support I've received at TELIN. Davy and Philippe, you were always ready to help and fix anything that could go wrong. I can't imagine how much our efficiency would drop without you to help us out. Annette, Alice, Patrick and Sylvia, it is amazing how you form a welcome oasis of calmness in the sometimes chaotic surroundings of TELIN. You always made me feel welcome at TELIN. I also would like to thank my fellow researchers at TELIN for all the interesting discussions, feedback, brainstorms and comments. And of course I would like to thank my (ex) office colleagues Jan, Simon, Michiel, Dirk, Wenzhi, Daniele, Junzhi and Filip for creating a pleasant atmosphere at work. Thanks!

Ghent, June 2014
Jonas De Vylder

Samenvatting

Visuele waarnemingen zijn altijd cruciaal geweest voor het verbeteren van het menselijk inzicht in complexe biologische processen. Met de ontwikkeling van computers zijn vele visuele waarnemingstaken (gedeeltelijk) geautomatiseerd. Daarnaast zijn er recent ook veel nieuwe geavanceerde beeldvormingstechnieken ontwikkeld. Deze nieuwe technieken maken het mogelijk om data van macroscopische tot moleculaire schaal te verzamelen, alsook het visualiseren van objecten door de tijd heen: observaties over verschillende weken heen tot metingen met microseconden tussen. Deze nieuwe beeldmodaliteiten leiden tot een tsunami aan data. De combinatie van de digitalisering en de ontwikkeling van nieuwe beeldvormingstechnieken heeft een renaissance in de biologie teweeg gebracht. Golven van celdelingen tijdens de ontwikkeling van een foetus, high-throughput analyse van fenotypes van planten, het in kaart brengen van cognitieve processen op cellulaire schaal, pre-symptomatische detectie van stress op planten, enz. Dit zijn slechts een paar van de baanbrekende onderzoeksthema's die momenteel worden onderzocht gebruikmakend van gedigitaliseerde biologische beeldvorming.

De analyse van deze grote en complexe datasets is ondertussen uitgegroeid tot een moeilijke, moeizame en tijdrovende taak voor biologen. Experts moeten relevante informatie uit grote datasets visualiseren en analyseren. Dit moet bovendien gebeuren voor grote en uitdagende datasets zoals time-lapse sequenties, 3D-beelden, multispectrale beelden, enz. Zulks een manuele analyse kan variëren van het zoeken naar een cel die een bepaald interessant proces ondergaat, bijvoorbeeld mitose of apoptose, of het handmatig tellen van interessante objecten, bijvoorbeeld het aantal celkernen of schimmelinfecties, tot de arbeidsintensieve taak van het manueel omlijnen van bepaalde objecten, zoals het omlijnen van een cel om de grootte te kunnen opmeten, of om bepaalde fenotypes te analyseren. Er is een grote nood aan methoden die de drastisch toenemende hoeveelheid biologische beeldmateriaal helpen te verwerken, dit zowel voor dagelijks laboratoriumwerk als voor high-throughput screening. Deze methoden moeten in hoge mate automatisch, computationeel efficiënt en nauwkeurig zijn. Hoewel er reeds veel computervisie-algoritmes ontwikkeld zijn voor toepassingen in andere onderzoeksdomeinen, zijn er slechts weinig algoritmes geschikt voor de uitdagingen die gepaard gaan met biologische beelden. Daarom moeten ofwel bestaande algor-

itmes significant aangepast worden, ofwel moeten er nieuwe specifieke methoden ontwikkeld worden.

In dit proefschrift bespreken we methoden om te helpen bij de automatische analyse van biologische afbeeldingen. De ontwikkelde methoden kunnen worden onderverdeeld in twee belangrijke groepen: 1) algoritmes voor automatische detectie en omlijning van relevante objecten, en 2) algoritmes voor de analyse van hun vorm.

In het kader van de automatische omlijning van objecten, ook wel segmentatie genoemd, hebben we het populaire actievecontour-raamwerk op verschillende manieren verbeterd. Actievecontour-segmentatie is een techniek waarbij een contour wordt verplaatst en vervormd in het ruimtelijke domein opdat een bijbehorende energiefunctie zijn minimum bereikt. De energie moet minimaal zijn als de contour het relevante object omlijnt. De meeste van onze voorgestelde verbeteringen passen de energie-optimalisatie op een zodanige manier aan dat deze technieken nu ook gebruikt kunnen worden voor een aantal belangrijke biologische taken. Deze biologische taken omvatten het volgen van meerdere objecten in lage-temporele-resolutie datasets, het nauwkeurig segmenteren van objecten in laagcontrast-infraroodbeelden, computationeel efficiënte segmentatie voor zeer hoge-ruimtelijke-resolutie-beelden zoals elektronenmicroscopie-beelden, segmentatie van objecten onder niet-uniforme belichting, enz. Wij stellen ook een nieuwe benadering van actievecontour-segmentatie voor: in plaats van het minimaliseren van een willekeurige energiefunctie, stellen wij voor om de probabiliteit dat een contour een object omlijnt te maximaliseren. We tonen aandat onze probabilistische actieve contouren een veralgemening zijn van enkele van de meest populaire actieve contouren. Deze Bayesiaanse aanpak heeft het voordeel dat alle parameters een statistische interpretatie hebben, in plaats van slechts ad hoc gewichten te zijn.

Als speciale use case voor beeldanalyse onderzochten we de automatische detectie en segmentatie van individuele celkernen in fluorescentiemicroscopie-beelden. We bestudeerden zowel methoden die gebruik maken van gedetecteerde randen als methoden die gebruik maken van de intensiteit. Voor de randgebaseerdemethoden, hebben we ons vooral gericht op de ontwikkeling van snelle en computationeel efficiënte segmentatiemethoden, terwijl voor de intensiteitsgebaseerdemethoden de focus meer op de ontwikkeling van accurate segmentatie methoden is gelegd. Hierbij werden vooral oplossingen gezocht voor de segmentatie van uitdagende 2D en 3D-datasets die naast celkernen ook irrelevante objecten bevatten, of voor clusters die een hoge densiteit aan celkernen bevatten. Vooral algoritmes die voorkennis over de vorm van de celkernen benutten leiden tot nauwkeurige celkerndetectie en -segmentatie. Uit grondige validatie, die zowel op gesimuleerde als op echte datasets is uitgevoerd, blijkt dat de voorgestelde methoden robuust zijn tegen ruis,

wazigheid en zelfs correlatie tussen verschillende kleurkanalen.

Binnen de groep van vormgerelateerde algoritmes, stellen wij twee verschillende technieken voor: 1) een algoritme om de vervorming van een object doorheen de tijd te analyseren en 2) een nieuwe reeks van generieke vormdescriptoren.

Voor het bestuderen van de vervorming van een object, stellen we een techniek voor die eerst de corresponderende contourpunten zoekt, m.a.w. we zoeken voor ieder punt van de contour op tijdstip A een corresponderend punt op tijdstip B. Op basis van deze correspondentie van de contourpunten wordt de vervorming van het inwendige van het object berekend. In tegenstelling tot vele methoden uit de literatuur, maakt deze aanpak geen assumpties omtrent de te bestuderen vormen. De voorgestelde methode werd gebruikt om de vervorming van aberrante celkernen te bestuderen. Ons algoritme kan de beweging van subnucleaire deeltjes onderscheiden van globale nucleaire vervormingen. Dit verbetert de mobiliteitsanalyses van genoom-loci en bepaalde eiwitstructuren aanzienlijk, zelfs binnen sterk vervormde celkernen. Deze sterke vervormingen komen dikwijls voor bij specifieke ziektes, zoals laminopathies of kanker.

Tot slot presenteren we een nieuwe groep vormdescriptoren. Vormdescriptoren kunnen gebruikt worden voor vele verschillende taken: om een bepaald fenotype te beschrijven, om gebruikt te worden als voorkennis bij segmentering, om gelijkaardige vormen terug te vinden in een databank, om bepaalde vormen automatisch te classificeren, enz. Veel state-of-the-art vormdescriptoren projecteren een parametrisatie van de contour van een vorm op een orthogonale basis. In dit proefschrift stellen we een soortgelijke aanpak voor, maar in plaats van een vaste parametrisatie te gebruiken, verbeteren we de vormbeschrijving door deze parametrisatie te optimaliseren. Wij combineren deze aanpak met de populaire Fourier descriptoren. Hierdoor kunnen we accuraat vormen beschrijven ongeacht huun locatie of oriëntatie. Onze voorgestelde vormdescriptoren realiseren een significante verbetering op het vlak van vormreconstructie, vormclassificatie en het vinden van gelijkaardige vormen in een databank.

In de loop van dit doctoraatsonderzoek zijn er 40 wetenschappelijke artikelen gepubliceerd als een gevolg van het onderzoek gebundeld in dit proefschrift. Van deze artikelen werden er vijf gepubliceerd in internationale wetenschappelijke tijdschriften geïndexeerd in de ISI Web of Knowledge-databank. Daarnaast lanceerden we RosetteTracker, een freeware tool voor de automatische analyse van Arabidopsis rozetten in topviewbeelden.

Summary

Visual observations have always been part of improving the human understanding of complex biological processes. With the evolution of computers, many visual observation tasks have been (partially) automated. Furthermore, many new sophisticated imaging modalities have been developed recently, which allow data acquisition from macroscopic to molecular level, and on a time-scale ranging from weeks to microseconds. These new image modalities result in a tsunami of data. The combination of digitalisation with the development of new imaging techniques is driving a renaissance in biology. Cell division waves during embryogenesis; high throughput plant phenotyping; a mapping of cognitive processes on cellular level; and presymptomatic detection of stress on plants are just a few of the breakthrough research topics that are currently being investigated by the use of digitised bio-imaging.

However, the analysis of these large and complex datasets has become a complex, tedious, and time-consuming task for biologists. Experts have to visualise and manually extract relevant information from large datasets, such as time-lapse sequences, 3D images, multispectral images, etc. This manual analysis can range from searching a micrograph for a cell undergoing a certain process of interest, e.g. mitosis or apoptosis; to manually counting objects of interest, e.g. count the number of cell nuclei or mould infections; to very labour intensive delineating objects of interest, e.g. to measure cell size, or to analyse plant phenotypes. From daily lab work to high-throughput settings, there is a great need for methods that process the ever increasing amount of biological image data. These methods should be to a large degree automatic, computational efficient, and accurate. Although many computer vision algorithms for automatic image analysis exist in other fields, most of them are ill suited for the challenges arising with bio-images. Instead, algorithms have to be adjusted or new dedicated methods have to be devised.

In this dissertation, we develop methods to aid in the automatic analysis of biological images. The proposed methods can be divided into two large groups:
1) algorithms for the automatic detection and delineation of objects of interest, 2) methods for the analysis of their shape.

In the context of automatic delineation of objects, also called segmentation, we have improved the popular active contour framework in various ways. Active contours consist of a type of algorithms where contour is displaced and deformed in the spatial domain until an associated energy functional reaches its minimum. The energy should be minimal if the contour delineates the object of interest. Most of our proposed improvements adjust the energy optimisation such that specific biological task can be handled with the active contour framework: tracking multiple objects in low temporal resolution datasets, accurate segmentation of low contrast objects in infrared images, computationally efficient segmentation for very high spatial resolution images such as electron microscopy images, segmentation of objects under non-uniform lighting conditions, etc. We also propose a new approach to active contour segmentation: instead of minimising an arbitrary energy function, we propose to maximise the probability that a contour delineates an object. We prove that some of the most popular active contour algorithms are a specific case of our probabilistic active contours. This Bayesian approach has the benefit that all parameters have a statistic interpretation, instead of nearly being ad hoc weights.

As a special use case for image analysis, we investigated the automatic detection and segmentation of individual cell nuclei in fluorescence microscopic images. Both methods relying on edges and on intensity as salient feature where studied. For the edge-based methods, we focussed on fast and computational efficient segmentation methods, whereas for the intensity-based methods the main goal was to develop accurate segmentation methods, even for challenging 2D and 3D datasets containing clutter and densely clustered nuclei. Especially algorithms that exploit prior knowledge on the shape of the nuclei result in accurate segmentation and detection results. Extensive validation both on simulated and real dataset show that the proposed methods are robust against noise, blur and even bleed through.

Within the group of shape related algorithms, we proposed two different techniques: 1) an algorithm to analyse the deformation of a shape over time, 2) a new set of generic shape descriptors.

For the evaluation of shape deformation over time, we propose a technique that first finds correspondence between the contour points of the shape at two consecutive time frames. Based on this matching of contour points, the deformation of the interior of the shape is calculated. In contrast to many methods from literature, this approach does not make any assumptions on the shapes under investigation. The proposed method is used to study the deformation of aberrant cell nuclei. Our algorithm enables the separation of subnuclear particle motion from global nuclear deformations. This significantly improves mobility measurements of genome loci as well as nuclear protein bodies, even within strongly deforming nuclei associated with different pathologies, such as laminopathies or cancer.

Finally, we present a new set of shape descriptors. Shape descriptors can be used for many different tasks: to describe a certain phenotype, to be used in segmentation as prior knowledge on the object of interest, to retrieve similar shapes from a database, to cluster or classify shapes, etc. Many state-of-the-art shape descriptors project a parameterisation of the shapes contour onto an orthogonal basis. In this dissertation, we propose a similar approach, but instead of using a fixed parameterisation, we improve the shape reconstruction by optimising the contour parameterisation. We combine this approach with the popular Fourier descriptors. This enables accurate shape descriptors that are invariant to a number of rigid transformations, e.g. translation or rotation. Therefore, our proposed shape descriptors significantly improve shape reconstruction, shape classification and the retrieval of similar shapes in a database.

During the course of this Ph.d. research, 40 scientific papers have been published as a result of the work synthesised in this dissertation. Out of these, five where published in international peer-reviewed journals indexed in the ISI web of knowledge database. Furthermore, we launched RosetteTracker, a freeware tool for the automatic analysis of Arabidopsis rosettes in topview images.

List of Acronyms

ADMM	Alternating Direction Method of Multipliers
AITD	Autoimmune Thyroid Diseases
AUC	Area Under the Curve
BVF	Boundary Vector Field
CAD	Computer Aided Diagnosis
CLEM	Controlled Light Exposure Microscopy
CPM	Charged Particle Model
CPU	Central Processing Unit
CSS	Curvature Scale Space
DAPI	4´,6-diamidino-2-phenylindoled
DIC	Differential Interference Contrast
DNA	Deoxyribonucleic acid
DTW	Dynamic Time Warping
EM	Electron Microscopy
ER	Endoplasmic Reticulum
FD	Fourier Descriptors
FFT	Fast Fourier Transform
FISH	Fluorescence in Situ Hybridization
FLIM	Fluorescence Lifetime Imaging Microscopy
FM	False Matches
FN	False Negatives
FP	False Positives
FPR	False Positive rate
FRAP	Fluorescence Recovery After Photobleaching
FRET	Fluorescence Resonance Energy Transfer
GFP	Green Fluorescence Protein
GT	Ground Truth
GVF	Gradient Vector Flow
GPU	Graphics Processing Unit
HD	Hausdorff Distance
ICP	Iterative Closes Point
ICR	Image Content Retrieval

IFD	Improved Fourier Descriptors
IFFT	Inverse Fast Fourier Transform
IGVF	Intensity Gradient Vector Flow
IR	Infrared
LSFM	Light Sheet Fluorescence Microscopy
MAP	Maximum A Posteriori
MHD	Modified Hausdorff Distance
ML	Maximum Likelihood
MSD	Mean Squared Distance
PAM	Pulse Amplitude Modulated
PBMC	Peripheral Blood Mononuclear Cell
PCA	Principal Component Analysis
PTM	particle-tracking microrheology
RANSAC	Random Sample Consensus
RGR	Relative Growth Rate
RMSE	Root Mean Squared Error
ROC	Receiver Operating Characteristic
SCGVF	Shape Constrained Gradient Vector Flow
SEM	Scanning Electron Microscope
TEM	Transmission Electron Microscope
TP	True Positive
TPR	True Positive Rate
TV	Total Variation
VZC	Virtual Electric Field
VFC	Vector Field Convolution
VIS	Colour within the visible spectrum
ZM	Zernike moments

Table of Contents

1 Introduction **1**

1.1 Problem statement and topical outline 1

1.2 Contributions and list of publications 2

1.3 Organisation of this dissertation 5

I Background **9**

2 Bio-imaging **11**

2.1 Macroscopic imaging 13

2.1.1 Current methods 13

2.1.1.1 High throughput imaging 13

2.1.1.2 Multispectral imaging 14

2.1.2 Recent trends and future prospects 15

2.1.2.1 3D imaging 15

2.1.2.2 New image-modalities 16

2.1.2.3 In the field imaging 16

2.2 Microscopic imaging 17

2.2.1 Current methods 17

2.2.1.1 Bright-field microscopy 17

2.2.1.2 Fluorescence microscopy 18

2.2.1.3 Confocal microscopy 21

2.2.1.4 Optical contrast enhanced imaging 23

2.2.1.5 Electron microscopy 24
2.2.2 Recent trends and future prospects 25
2.2.2.1 Increasing spatial resolution 25
2.2.2.2 Increasing temporal resolution 27
2.2.2.3 Digital image restoration 29
2.2.2.4 Multimodal imaging 29
2.3 Conclusion 31

3 Image analysis **33**
3.1 Correlation based analysis 34
3.2 Edge based segmentation 35
3.2.1 Edge detectors 36
3.2.2 Tracing 36
3.2.2.1 Line tracing 37
3.2.2.2 Network tracing 39
3.2.3 Watershed 42
3.3 Intensity based segmentation 43
3.3.1 Blob detection 43
3.3.1.1 Detection based on linear filters 43
3.3.1.2 Detection based on morphological operators . . 44
3.3.2 Thresholding methods 48
3.3.3 Clustering based methods 50
3.4 Model based segmentation 50
3.4.1 Implicit models 51
3.4.2 Constrained methods 51
3.4.3 Energy optimisation methods 52
3.5 Experimental results 52
3.5.1 Use case 1: cell segmentation in DIC images 52
3.5.1.1 Background 53

3.5.1.2 Segmentation of pavement and guard cells . . . 54
3.5.2 Use case 2: A. thaliana rosette phenotyping 56
3.5.2.1 Background 56
3.5.2.2 Quantification of rosette growth using VIS images 58
3.5.2.3 Analysing a composite time lapse sequence of fluorescence and thermal (infrared) images . . 59
3.6 Conclusion 62

II Segmentation & tracking 63

4 Parametric active contours 65
4.1 Background 66
4.1.1 Regularisation 67
4.1.1.1 Generic internal energies 67
4.1.1.2 Higher order internal energies 69
4.1.1.3 Shape based internal energies 70
4.1.2 Data-fit 70
4.1.2.1 Data-fit with robustness to local optima 71
4.2 Intensity based data-fit 73
4.2.1 A joint edge and intensity based force field 74
4.2.2 Experimental results 77
4.3 Optimisation in time lapse sequences 80
4.3.1 Two phase optimisation 82
4.3.2 Experimental results 83
4.4 Optimisation using shape constraints 86
4.4.1 Shape constrained force 86
4.4.2 Experimental results 88
4.5 Optimisation using force propagation 89
4.5.1 Two pass propagation 90

4.5.2 Parameter settings 93
4.5.3 Directional force 94
4.5.4 Experimental results 96
4.6 Probabilistic active contours 99
4.6.1 Internal probability 100
4.6.2 External probability 100
4.6.3 Optimisation 101
4.6.4 Experimental results 102
4.7 Conclusion 107

5 Geometric active contours 109

5.1 Background 110
5.1.1 Mumford-Shah piecewise constant active contours 110
5.1.2 Globally optimal active contours 113
5.2 Probabilistic geometric active contours 115
5.2.1 The image model 116
5.2.1.1 Piecewise constant Mumford-Shah 116
5.2.1.2 General stationary image models 118
5.2.1.3 Non-stationary image models 121
5.2.2 A probabilistic prior 129
5.2.2.1 Perimeter & area 130
5.2.2.2 Markov random fields 130
5.2.2.3 Shape descriptors 133
5.3 Optimisation 134
5.3.1 Gradient descent 134
5.3.2 Newton-Raphson 135
5.3.3 Line search 136
5.3.4 Dual formulations 136
5.3.5 Split alternating direction method of multipliers 138

5.3.6 Other methods 139
5.3.7 Experimental results 140
5.3.7.1 Convergence 141
5.3.7.2 Scalability 143
5.3.7.3 Robustness to noise 144
5.4 Conclusion 145

6 Detection and segmentation of cell nuclei 147

6.1 Edge based nuclei detection 149
6.1.1 Use case: fibroblast nuclei 150
6.1.2 Use case: Hybridoma 152
6.1.2.1 Edge propagation 153
6.1.2.2 Nuclei segmentation 157
6.1.2.3 Experimental results 158
6.2 Intensity based nuclei detection 160
6.2.1 Use case: embryonic cell nuclei 161
6.2.1.1 Local intensity 161
6.2.1.2 Experimental results 162
6.2.2 Use case: Neuronal cell nuclei 164
6.2.2.1 Contrast enhancement 167
6.2.2.2 Nuclei detection 170
6.2.2.3 Error correction 174
6.2.2.4 3D Merging and Splitting 178
6.2.2.5 Experimental results 179
6.2.3 Use case: peripheral blood mononuclear cell nuclei 184
6.2.3.1 Joint shape and intensity constraints 186
6.2.3.2 Nuclei detection 189
6.2.3.3 Experimental results 192
6.3 Conclusion 199

III Shape & deformation 201

7 Cell nuclei mapping 203

7.1 Segmentation . . . 205
7.2 Interphase mapping . . . 205
7.3 Nuclear mapping of pathological cells . . . 206
7.3.1 Contour matching . . . 208
7.3.2 Cell mapping . . . 212
7.3.3 Telomere tracking . . . 214
7.3.4 Experimental results . . . 214
7.3.5 Discussion . . . 222
7.4 Conclusion . . . 223

8 Shape description 225

8.1 Scalar morphometrics . . . 228
8.2 Region-based descriptors . . . 229
8.2.1 Geometric moments . . . 231
8.2.2 Hu-Li moments . . . 232
8.2.3 Zernike moments . . . 233
8.2.3.1 Invariance . . . 236
8.3 Contour-based descriptors . . . 236
8.3.1 Shape description using an orthogonal basis . . . 237
8.3.1.1 Invariance . . . 238
8.3.2 Shape description using a warped orthogonal basis . . . 240
8.3.2.1 Optimisation . . . 241
8.3.3 Shape description using a frame . . . 243
8.4 Experimental results . . . 243
8.4.1 Shape representation . . . 244
8.4.2 Shape classification . . . 249
8.4.3 Image content retrieval . . . 249
8.5 Conclusion . . . 253

9 Conclusion **255**

9.1 Review of our contributions 255

9.2 Directions for future research 257

A Data sets **261**

A.1 Fibroblast cells 261

A.1.1 Cell culture and transfection 261

A.1.2 Image acquisition 262

A.2 Hybridoma cells 262

A.2.1 Cell culture and transfection 262

A.2.2 Image acquisition 262

A.3 Embryonic cells 263

A.3.1 Cell culture and transfection 263

A.3.2 Image acquisition 263

A.4 Peripheral blood mononuclear cells 264

A.4.1 Cell culture and transfection 264

A.4.2 Image acquisition 264

A.5 VIS images A. thaliana 264

A.5.1 Ecotypes and growth 264

A.5.2 Image acquisition 264

A.6 DIC images A. thaliana 265

A.6.1 Ecotypes and growth 265

A.6.2 Image acquisition 265

A.7 Chlorophyll fluorescence and IR images A. thaliana 265

A.7.1 Ecotypes and growth 265

A.7.2 Image acquisition 265

A.8 Chlorophyll fluorescence and IR images of sugar beet seedlings . 266

A.8.1 Ecotypes and growth 266

A.8.2 Image acquisition 266

A.9 Digital herbarium 266

A.9.1 Ecotypes and growth 266

A.9.2 Image acquisition 266

B Quality metrics **269**

B.1 Detection metrics 269

B.2 Segmentation metrics 270

B.3 Tracking metrics 270

B.4 Shape difference metrics 271

B.5 Classification metrics 271

C Curriculum Vitae **273**

1 Introduction

In this dissertation, we examine and develop specific computer vision algorithms for biological image analysis. The proposed work focuses both on finding practical solutions for biological imaging problems and on theoretical improvements of computer vision algorithms.

1.1 Problem statement and topical outline

The latest United Nations projections indicate that world population will reach 10 billion people around 2062 [UN 12]. At the same time, life expectancy forecasts predict that people will generally grow significantly older, e.g. males born today in England and Wales can on average expect to live until 90, and for females this is 94, with 32% of males and 39% of females expected to reach 100 in 2112. So in concreto, more people will inhabit the world, and generally they will grow older as well. This population growth will put a huge strain on resources and on society in general. There is still a long way to go to reach the point that all of these people will have the means to survive, and moreover that they can live in reasonable circumstances and quality. For example in order to avoid famine, crop yield should increase 2.4% each year to sustain the expected population. A different example relates to healthcare: it would be desirable to spend the life expectancy gain in good health, instead of living these extra years in pain or oblivious of our surrounding due to dementia or Alzheimer disease.

To cope with the challenges ahead, biology will have to make significant progress to fully understand all cellular and genetic processes. Even with the development of many chemical tests, images remain one of the most important ways to observe and analyse biological phenomena. In contrast to chemical tests, images contain spatial information, which can be highly valuable for biological experiments. However, acquisition and analysis of time-dependent cellular processes is very laborious. Hence there is a great need for automated methods to speed up the analysis of bio-images.

Throughout this thesis we will mainly focus on algorithms that can help biologists in their studies. The proposed algorithms enable efficient and automated analysis of images. Many of these algorithms belong to the class of image segmentation techniques, which are well studied within computer vision. Some examples are algorithms that detect individual cell nuclei in a microscopy image that visualises large groups of clustered nuclei, or algorithms that delineate full rosettes of Arabidopsis plants in infrared images or in time lapse sequences.

Next to image segmentation applications, we also consider shape analysis methods. These can be used once the objects of interest have been detected. The proposed methods allow evaluating and quantifying the difference between two shapes. This is both useful for the analysis of the morphological evolution of a shape over time, and for the comparison of two different object, e.g. to retrieve objects with similar shape from a database.

1.2 Contributions and list of publications

This research has resulted in several contributions. The key novelties that will be discussed in detail in this dissertation are:

- We propose several methods for tracking and segmenting multiple objects in time lapse sequences. These methods result in accurate tracking results by using prior knowledge on the intensity, shape, and/or motion of the objects of interest. This work has been done in collaboration with Dominique Van Der Straeten, and Laury Chaerle (Laboratory of Functional Plant Biology, Ghent University) and with Daniel Ochoa Donoso ((Escuela Superior Politécnica del Litoral, ESPOL University)).
- In order to segment cell wall networks, we propose a novel segmentation algorithm. The proposed method is a generalised tracing algorithm that is able to track bifurcating lines. This makes it suitable to segment network structures. The method is applicable on multiple microscopic modalities such as fluorescence, but also for images captured using a non invasive DIC

microscope. This work has been done in collaboration with Stijn Dhondt and Dirk Inzé (Department of Plant Systems Biology, Vlaams instituut voor biotechnologie, Ghent University).

- We introduce Rosette Tracker, a new freeware image analysis tool for evaluation of plant shoot phenotypes (http://telin.ugent.be/~jdvylder/RosetteTracker). This tool is not constrained to one specific monitoring system and can be adapted to different low budget imaging setups and requires minimal user input. In contrast to state-of-the-art automated plant monitoring tools, Rosette Tracker allows simultaneously quantifying plant growth, photosynthesis and leaf temperature related parameters through the analysis of VIS (common colour image in the visible spectra), chlorophyll fluorescence and/or thermal infrared time lapse sequences. Rosette Tracker facilitates gaining rapid understanding of Arabidopsis thaliana genotype effects. With over 1100 downloads of Rosette Tracker, the demand of the phytologist community for such a tool becomes apparent. This work has been done in collaboration with Dominique Van Der Straeten, Filip Vandenbussche and Yuming Hu (Laboratory of Functional Plant Biology, Ghent University).

- Active contours or snakes are widely used for segmentation and tracking. Although the active contour framework has proven its merit, there is a clear lack of a theoretical framework that enables an optimal setting of parameters and energy terms. In this work, we propose such a theoretical framework based on Bayesian theory. This probabilistic approach results in a generalisation of state-of-the art active contour segmentation. In the proposed framework all parameters have a statistical interpretation, thus avoiding ad hoc parameter settings.

- We propose a new class of shape descriptors. By optimising the scanning speed of the contour, we achieve better shape descriptors. This approach results in more accurate shape approximation, better shape classification and better shape retrieval.

- Many segmentation and tracking algorithms require the minimisation of an energy function. Within this dissertation, we investigate and evaluate several strategies for the optimisation of both non-convex and convex energy functions.

- We propose a novel and robust cell nuclei detection method based on the active contour framework. Improvement over conventional approaches is achieved by using prior knowledge of the nucleus shape in order to better

detect individual nuclei. This prior knowledge is defined using a dictionary based approach, which can be formulated as the optimisation of a convex energy function. This work has been done in collaboration with Mado Vandewoestyne, Trees Lepez and Dieter Deforce (Laboratory for Pharmaceutical Biotechnology, Ghent University) and in collaboration with Bruce McNaughton and his team (Polaris Brain Dynamics research group, The University of Lethbridge, Canada).

- In order to study the dynamics of genome regions, it is necessary to remove perturbations due to motion and deformation of the nucleus, i.e. the container holding the genome. While several models have been proposed to cope with the transformations corresponding to nuclear dynamics of healthy cells, none of these models are suited for pathological cells. In this thesis we propose a method that does not constrain the possible nuclear shapes nor on the possible deformations, which makes it well suited for the analysis of deformations occurring in pathological nuclei. This work has been done in collaboration with Winnok De Vos (Cell Systems & Imaging, Ghent University).

So far, our work resulted in 3 published international peer reviewed publications (as first author), 2 international journal publications as second author, and one publication in a national journal. One journal paper as second author is submitted and 4 international journal publication as first author are being prepared. 34 papers have been published in the proceedings of international conferences (28 as first author). A small selection of the key publications published during this research is given below, for a complete list of our publications we refer to Appendix C.

Shape related publications:

- J. De Vylder, W. H. De Vos, E. M. Manders & W. Philips. 2-D Mapping of Strongly Deformable Cell Nuclei, Based on Contour Matching". *Cytometry Part A*, vol. 79A, no. 7, pages 580-588, 2011 [impact factor = 3.711].

- J. De Vylder & W. Philips. "2-D Shape Representation Using Improved Fourier Descriptors". In *16th IEEE International Conference on Image Processing (ICIP)*, pages 397-400, 2009.

Active contour related papers:

- J. De Vylder & W. Philips. "A Computational Efficient External Energy for Active Contour Segmentation Using Edge Propagation". In *17th IEEE International Conference on Image Processing (ICIP)*, pages 661-664, 2010.

- J. De Vylder, D. Van Haerenborgh, J. Aelterman & W. Philips. "A Probabilistic Interpretation of Geometric Active Contour Segmentation". In *European Signal Processing Conference (EUSIPCO)*, 2014.
- J. De Vylder, D. Ochoa, W. Philips, L. Chaerle & D. Van Der Straeten. "Tracking Multiple Objects Using Moving Snakes". In *16th International Conference on Digital Signal Processing (DSP)*, pages 966–971, 2009
- J. De Vylder, J. Aelterman & W. Philips. "A comparative study on the optimisation of global optimum active contours". In *5th IEEE International Conference on Signal Processing and Communication Systems*, December 2011.

Application oriented papers:

- J. De Vylder, J. Aelterman, M. Vandewoestyne, T. Lepez, K. Douterloigne, D. Deforce & W. Philips. "A Novel Dictionary Based Computer Vision Method for the Detection of Cell Nuclei." *PLOS ONE*, vol. 8, no. e54068, 2013 [impact factor = 3.730].
- J. De Vylder, F. Vandenbussche, Y. Hu, W.Philips & D.Van Der Straeten. "Rosette Tracker: an open source image analysis tool for automatic quantification of genotype effects". *Plant Physiology*, 2012 [impact factor = 6.535].
- J. De Vylder & W. Philips. "Computational efficient segmentation of cell nuclei in 2D and 3D fluorescent micrographs." In *Imaging, Manipulation, and Analysis of Biomolecules, Cells, and Tissues IX, proceedings of SPIE,* Vol. 7902, 2011.
- J. Aelterman, B. Goossens, J. De Vylder, A. Pižurica & W. Philips. "Computationally efficient Locally Adaptive Demosaicing of Color Filter Array Images using the Dual-Tree Complex Wavelet Packet Transform." *PLOS ONE*, vol. 8, no. e61846, 2013 [impact factor = 3.730].
- D. Babin, A. Pižurica, J. De Vylder, E. Vansteenkiste & W. Philips. "Brain blood vessel segmentation using line-shaped profiles". *Physics in medicine and biology*, vol. 58, no. 22, pages 8041–8061, 2013 [impact factor = 2.701].

1.3 Organisation of this dissertation

This thesis is organised in three large parts. The first parts gives an introduction and some background on bio-imaging and on automated analysis of bio-images.

The second part elaborates on specific methods used for segmentation and tracking of objects. The third part discusses the concepts of shape description and shape deformation.

Part 1: Background

In Chapter 2 we review different image modalities used within biology. We discuss macroscopic imaging using high throughput systems using different cameras such as colour cameras, depth cameras, infrared cameras, etc. We also present an overview of microscopy systems. We describe different optical microscopes such as brightfield, confocal, differential interference contrast (DIC), etc. We also give a short overview of electron microscopy.

In Chapter 3 we give some background on different strategies that can be used to automatically analyse images. We discuss the use of intensity and edges as main features. We also elaborate on more complex and more powerful model based methods. We show the potential of some of these methods by means of two use cases: the analysis of cells in DIC images, and the segmentation of Arabidopsis plants in colour and chlorophyll fluorescence time lapse sequences.

Part 2: Segmentation & tracking

In Chapter 4 we present the parametric active contour method. This is a model based segmentation method that exploits edges in an image. We propose several algorithms that fit within this framework, with a special emphasis on their usefulness for the analysis of time lapse sequences. We also propose a novel and efficient optimisation method for these type of active contours.

In Chapter 5 we present a new probabilistic framework for geometric active contour segmentation. We show that the proposed framework is generic and can incorporate many different image models and many different shape priors. We also investigate the use of different methods for the optimisation of the convex energies used in this segmentation framework.

In Chapter 6 we investigate several algorithms for the segmentation and detection of individual cell nuclei. We describe several challenges such as non-uniform lighting or the clustering of cell nuclei. We propose a number of segmentation methods, each suited for different challenges encountered in several specific cases.

Part 3: Shape & deformation

This part consists of chapters 7 and 8. Chapter 7 discusses the deformation analysis of pathological cell nuclei. These deformations hamper the study of subnuclear particle dynamics. In this chapter we propose a new method to remove all global motion induced by the motion and deformation of the nucleus.

In chapter 8 we present an overview of shape descriptors. We investigate a new method that optimises the scanning speed of a contour in function of better shape representation. We show that this approach results in shape descriptors that are better suited for shape classification, retrieval and representation.

Finally, in chapter 9 a general conclusion is given for our work and some directions for future research are discussed.

Part I

Background

2

Bio-imaging

Biology is intrinsically related to observing phenomena in nature. This can either be the case to form new models for specific biological processes, or to validate an already existing model such as a model of the effect of a specific gene on plant growth. While originally these observations where done real time and in-place, it already became clear a long time ago that being able to keep a visual record would be of great value. This would remove the spatial and temporal constraints, and thus enable the study and comparison of biological objects over time, or over long distances. In the 18th century these advantages where already apparent, even with the lack of easy access to imaging devices such as a camera. For example on the journey of the Beagle, John Gould sketched a large number of finches. This allowed comparing these finches, even if they were each endemic to different islands. Fig. 2.1 shows an example of four sketches of these finches of the Galapagos islands. These sketches, and the comparative studies that they enabled, formed the main inspiration for Darwin's theory of the origin of species. This bio-imaging avant la lettre shows clearly the added value of imaging for the study of living organisms. Over the years, bio-imaging has greatly evolved, e.g. allowing the visualisation of light at random wavelengths, to observe structures at nanometre scale, to measure 3D geometry, etc. In this chapter, we discuss the most important bio-imaging technologies, which we have grouped in microscopic and macroscopic imaging. For each group we discuss the current state-of-the-art and give a short overview of what the recent trends and potential future prospects are.

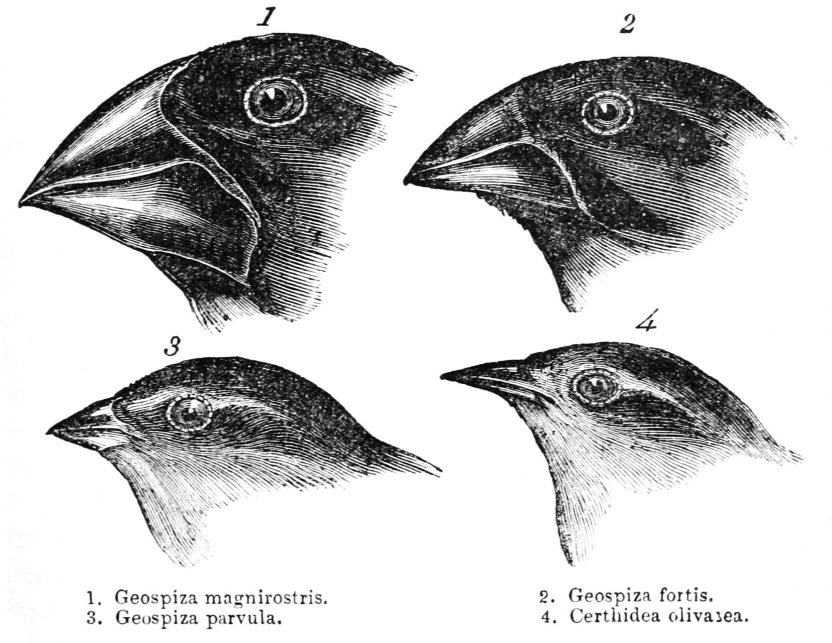


Figure 2.1: Sketches of finches at the Galapagos islands: Large ground finch, Medium ground finch Small tree finch, Green warbler-finch. From "Voyage of the Beagle".

2.1 Macroscopic imaging

2.1.1 Current methods

2.1.1.1 High throughput imaging

It has been demonstrated that the rosette area of a plant is directly proportional to its weight [Leister 99]. This observation opened the path for plant phenotyping based on images. The most important benefit of using images is that it generally results in a non-destructive and even non-obtrusive measurement, in contrast to more established measurements such as harvesting and weighing the biomass. This allows the analysis of plants over time. Given the added value of image based monitoring, numerous studies have been done using this approach. In order to meet the challenges of this new and promising field of plant phenotyping, new technologies such as greenhouse automation and plant imaging systems have been developed. These setups capture an image of up to thousands of plants a day. This large scale monitoring allows the simultaneous study of certain traits on different wild types and mutants.

Two groups of high throughput systems can be distinguished: systems that move each plant to an imaging chamber and systems where the cameras are moved from plant to plant. The benefit of the first system is that it requires less green house space, the plants can grow relatively close to each other since plants are imaged at a separate dedicated location. One of the criticisms on these type of systems is that the impact of the movement on the plant is insufficiently studied. Especially for images captured at wavelengths outside the visible wavelength range, caution is required. The reason for this is that the motion of plants in a still atmosphere will resemble the effect of wind on the plant. It is known that plants under the stress of wind will close there stomata, and as such will adjust their photosynthesis and temperature. These two features can be measured at wavelengths not necessarily within the visible range as will be discussed in section 2.1.1.2.

The setups that move the cameras have the benefit that plants remain stationary, so the impact of motion is not an issue for these systems. The downside is that the images are taken in place, which often results in images that are more challenging than images captured in a dedicated imaging chamber, e.g. a non-uniform background, different lighting conditions, clutter, etc. all hamper automatic image analysis. Fig. 2.2 shows two examples of high throughput imaging systems. On the left, the phenovision system at VIB is used to monitor corn plants. This system belongs to the first group, where the plants are move using conveyor-belts towards the imaging chamber (the white box in the image). The right part of Fig. 2.2 shows an example of automated imaging, where the plants remain in place, but where a

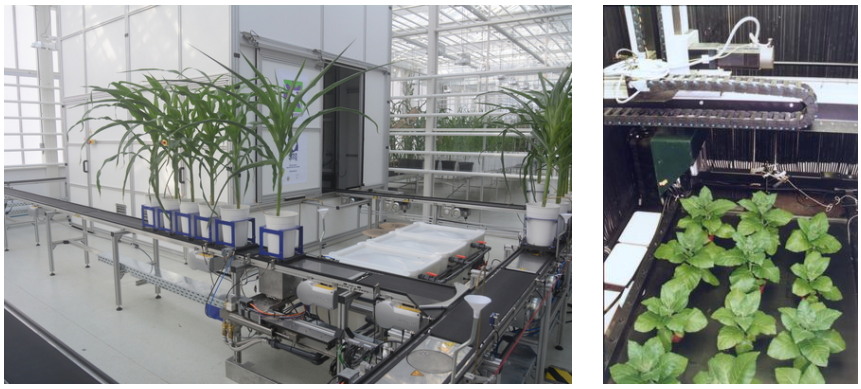


Figure 2.2: Two examples of automated plant imaging systems. Left, the phenovision robot at VIB, where plants are transported towards an imaging chamber (Image by courtesy of Stijn Dhondt, department of Plant Systems Biology, Ghent university - VIB). Right, the imaging system at the Laboratory of Functional Plant Biology, Ghent university. This system moves multiple cameras over the plants (Image by courtesy of Dominique Van Der Straeten, Laboratory of Functional Plant Biology, Faculty of Sciences, Ghent University).

set of cameras is moved using a robotic arm. This robotised setup is part of the Laboratory of Functional Plant Biology facility, Ghent university.

2.1.1.2 Multispectral imaging

By the time a plant displays visible symptoms of stress, it can already be adversely affected. Recent imaging techniques that capture wavelengths in a wider spectrum than the visible spectrum (390-700 nm) allow presymptomatic monitoring of changes in plant states. Especially fluorescence and infrared (IR) imaging have proven their merit by detecting stress-related changes in plant leaves. Their application for crop monitoring could permit to alleviate stress at an early stage, thus limiting yield loss by circumventing irreversible damage due to specific stress factors.

Fluorescence imaging can provide information on chlorophyll content without the need for pigment extraction. The chlorophyll content is correlated to photosynthesis within the leafs. Fluorescence can also be used for size estimations of green plants [Jansen 09].

IR images can be used to estimate leaf temperature. Temperature is dependent on evaporation and transpiration of the leaf, which is partly determined by stomatal opening. Hence, IR imaging can be used to monitor transpiration differences [Wang 03].

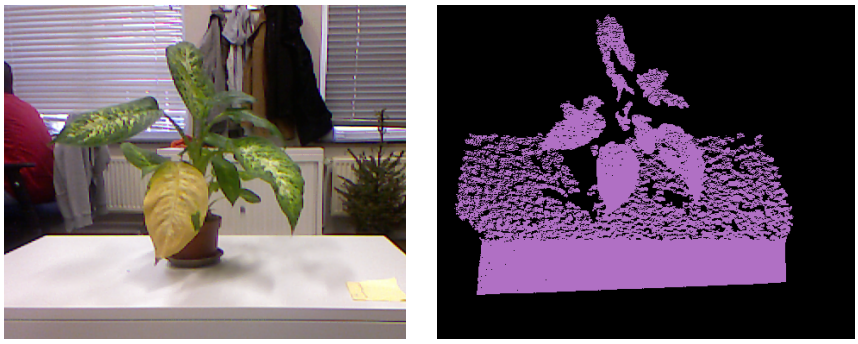


Figure 2.3: An example of 3D plant monitoring. Left, a regular VIS image of a plant, right a 3D point cloud from the same plant.

2.1.2 Recent trends and future prospects

2.1.2.1 3D imaging

Although several high throughput monitoring systems already capture a number of side view images in combination with a top view image, this mainly provides accurate estimates on biomass, but generally contains little information of the actual geometry of the plant. Recently, monitoring systems have been proposed that do actual 3D monitoring. This allows real 3D kinematic studies, e.g. to analyse plant growth in the close vicinity of other plants, or to measure leaf angle through a diurnal cycle [Dornbusch 12, Sirault 12].

There exist several methods to measure geometric information of an object, e.g. structured light, time of flight, structure from motion, etc. However, in the context of plant phenotyping the following two methods seem to be the most explored: Lidar and multi-camera 3D restoration. Lidar measures distance by illuminating a target with a laser and analysing the reflected light. This approach is used i.a. in the high throughput phenotyping system of Lemnatec. Multi-camera 3D restoration calculates the 3D geometry based on the 2D displacement of an object between two images. This is similar to how predators have depth perception due to the disparity of the view coming from both eyes.

Fig. 2.3 shows a proof of concept of 3D plant monitoring. The left part of the figure shows the plant using a regular VIS image. The right part shows a point cloud of the same plant. This point cloud is generated using off-the-shelf equipment, in this case, a kinect camera.

2.1.2.2 New image-modalities

In search of new opportunities to study plants, the added value of new imaging modalities are being tested. The use of medical imaging modalities is such an example. Plants are monitored using MRI, CT or μ-CT. These show the 3D structure of plants [Dhondt 12a, Pridmore 12, Schulz 12], even of plant parts that are normally not visible such as roots. Although the imaging modalities should still be optimised for this application, e.g. to reduce radiation in order to minimise the decrease of plant growth, the preliminary results look promising.

Another image modality that is explored is hyperspectral imaging. Instead of only monitoring three ranges of wavelengths, such as is custom in VIS imaging, or by extending this to a limited set of bands in multispectral imaging, the image is captured for a large range of wavelengths (over 200 bands). These bands, or a combination of some of these bands, can reveal plant water content, and by extension water stress and draught assessment [Weckler 04].

2.1.2.3 In the field imaging

The majority of plant phenotyping research focusses on the analysis of model organisms in lab conditions. Model organisms are species that are extensively studied in order to understand complex biological phenomena. Some of the most extensive studied model organisms are Drosophila melanogaster, Caenorhabditis elegans and Arabidopsis thaliana. Arabidopsis is a flowering plant related to cabbage and mustard, whose small stature and short generation time facilitates rapid genetic studies of plants. For an in-depth study of automated image analysis of model organisms, we refer to [Ochoa Donoso 11]. While laboratory analyses of key plant parts may complement direct phenotyping under field conditions, information captured throughout the plant life cycle in a carefully controlled environment, traits and candidate genes identified within controlled environments have generally not translated into gains in grain yield in the field [Araus 14]. Results from controlled environments are far removed from the heterogeneous environmental factors that plants will experience in the field and, therefore, are difficult to extrapolate to the field. “In the field” phenotyping is increasingly seen as key to achieving progress in crop improvement, although this brings its own challenges. Outdoor measurement systems have to deal with more variability compared to growth chamber phenotyping: variability of light, soil, external growth conditions, etc.

In the field monitoring is either done from a stationary portal crane system or by small drones such as provided by gatewing, Trimble. These drones can be programmed to periodically and automatically scan a test field using different image modalities.

2.2 Microscopic imaging

Biology already relies for a long time on observations made by microscopes. The use of lenses was known even in ancient China (~2000 BC), where a lens and a water-filled tube were used to visualize the unseen [Bardell 04]. Nevertheless, it wasn't until the sixteenth century that larger magnifications became theoretically possible by combining multiple lenses. Since the first report on the use of multiple lenses, microscopy has evolved on many levels: lenses have been combined to remove chromatic aberration, larger magnifications were achieved, the use of different light sources has been explored, the limitation of focal depth has been studied, i.e. in order to allow optical sectioning, etc. In this section, we will give an overview of the most relevant microscopic systems for this dissertation. Note that this is not an exhaustive list of all microscopic systems, for more in depth discussion on microscopy and its different setups, we refer to [Stelzer 98, Davidson 02, Vonesch 06, Tkaczyk 10].

2.2.1 Current methods

2.2.1.1 Bright-field microscopy

Dutch spectacles-makers Hans Jansen and his son Zacharias Jansen where the first to propose the use of a set of lenses in a tube. Their discovery around 1590 that the combination of multiple lenses resulted in theoretical higher magnifications than the magnification achieved by single lenses, resulted in a new research field that studies structures too small to be observed with the human eye. The first feat of arms of the so-called compound microscopes encompasses breakthrough observations of a fly's eye, of a flea, of cells, etc. Through history the understanding of optics and the manufacturing of lenses has greatly evolved thanks to pioneers such as Antony van Leeuwenhoek, Galileo Galilei, Ernst Abbe, Christian Huygens and many others. For a brief history on optical microscopy, we refer to [Bardell 04, Rooms 05b]

However, current bright-field microscopes still implement the basic concept of the original compound microscopes developed by the family Jansen in the late sixteenth century. Fig. 2.4 shows a schematic overview of the bright-field microscopic setup. A sample is illuminated from below by white light and observed from above. [1]

[1] Note that for simplicity that the lens system here is represented by a single lens, in practice this typically consists of a large set of lenses collected in a tube. One set of lenses is located at the objective of the microscope, and one lens at the eyepiece of the microscope. It is the combination of these two lens groups that results in high magnifications.

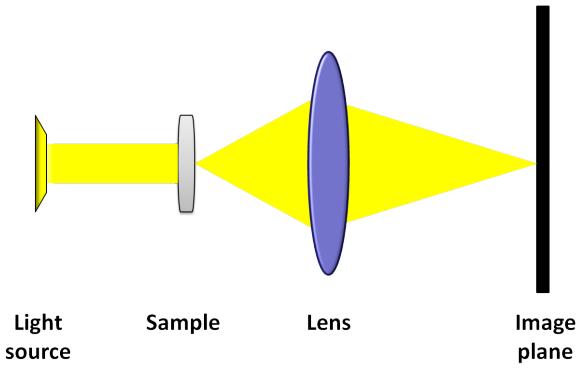


Figure 2.4: A schematic overview of bright-field microscopy. Note that the lens is a symbolic representation actually consisting of multiple lenses.

Since the light passes through the specimen, some part of it deviates as it encounters opaque parts of the specimen. This deviation of light is called diffraction. Fig. 2.5 shows a simple example of diffraction. In this example a planar wavefront with wavelength λ passes a barrier with two slits. This results in two circular wavefronts to the right of the barrier. These two wavefronts interfere with each other: at certain orientations these waves interfere destructively, whereas at other orientations the waves amplify each other (these orientations are called orders). The right part of Fig. 2.5 shows the actual diffraction pattern that would be captured at the objective plane. In contrast to this simple example with two slits, a microscopic specimen can be considered a complex grating. This concept was first used by Ernst Abbe to study the theory of image formation. He concluded that in order to accurately resolve the details of a specimen, that at least the 0^{th} and 1^{st} diffraction orders have to be captured by the objective lens. The more diffraction orders captured by the objective, the more accurate the image will represent the actual specimen, i.e. the higher the microscopic resolution.

The lens system focuses the diffracted light from the sample onto an image plane. Either this image plane corresponds to the location of a microscopists eye, or to the location of a sensor that captures a digital image. Contrast in the image is caused by absorption of some of the transmitted light in dense areas of the sample. For a more in-depth discussion on diffraction and microscopic image formation, we refer to [Davidson 02, Tkaczyk 10].

2.2.1.2 Fluorescence microscopy

Conventional optical microscopes are hampered by the lack of contrast in biological samples. By colouring the sample using specific dyes, this can be improved.

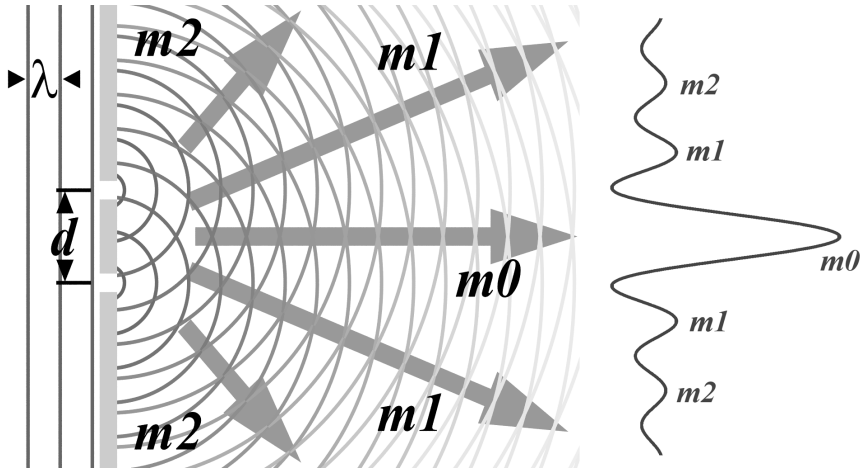


Figure 2.5: An illustration of the diffraction of a wavefront of wavelength λ passing two slits at a distance d of each other. The diffracted light passing the two splits interferes and amplifies at certain directions (called orders and represented by arrows). This results in the diffraction pattern on the right.

Histology is an example where specific cells or organelles are marked using specific dyes. Although these dyes increase the contrast, the use of white light that diffracts within the sample can cause interference: the observer receives light scattered from the region of interest as well as light scattered from surrounding regions.

This issue is solved by the use of fluorescent dyes. Especially the use of fluorescent proteins has revolutionised life science research. Fluorescent proteins belong to a class of proteins that share the unique property of being self-sufficient to form a visible wavelength chromophore. In fluorescence microscopy, a gene encoding an engineered fluorescent protein is introduced into living cells. The presence of a chromophore allows the visualization of the location and dynamics of the protein.

This is illustrated in Fig. 2.6. A broad spectrum light source is filtered such that only photons with a specific wavelength lights the sample. This light is absorbed by the fluorescent protein, which causes it to emit light at a different wavelength. Note that the excitation is also diffracted within the sample, so the light being focussed by the lens system is a combination of diffracted light originating form the light source and emitted light from the fluorophores. The emitted light from the fluorophores is separated from the rest by placing a colour filter on the optical path, thus blocking all light not emanating from the fluorescent proteins. These types of microscopes are called widefield microscopes. The use of fluorescent proteins has had such an important impact on understanding complex biological systems,

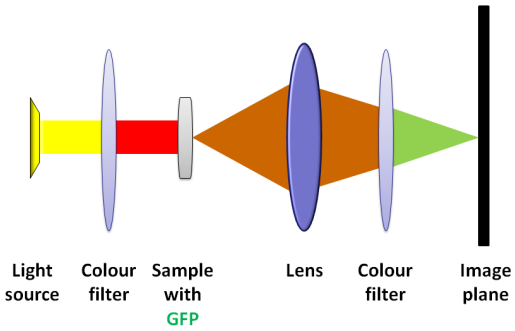


Figure 2.6: A schematic overview of fluorescence microscopy. The sample expresses the green fluorescence protein (GFP).

that Martin Chalfie, Osamu Shimomura, and Roger Y. Tsien were awarded the 2008 Nobel Prize in Chemistry for their discovery and development of the green fluorescent protein. Not only the visualisation of fluorescent proteins, but also the dynamics of fluorescence over time has improved the insight of internal cellular kinematic process, e.g. by fluorescent based techniques such as Fluorescence lifetime imaging (FLIM), Fluorescence recovery after photobleaching (FRAP), fluorescence resonance energy transfer (FRET), etc. [Vonesch 06]

In order to separate the excitation light from the laser and emission light from the fluorescent proteins, the fluorophores should emit at a different wavelength than the excitation wavelength. The difference between the excitation and emission wavelength is called the Stokes shift. A larger Stokes shift allows better separation of scattered illumination light and fluorescent light from the sample. Fig. 2.7 shows the emission and excitation spectra of six different common fluorescent proteins. Note that multiple fluorescent proteins can be used to dye a single sample. This allows visualising different proteins within the same sample. In such a case, it is beneficial to use fluorescent proteins with emission and excitation spectra that overlap as little as possible. Otherwise, fluorescence coming from one type of fluorescent protein, will be visible in the colour channel corresponding with the other fluorescent protein. This is called bleed-through and can severely hamper the analysis of the sample. A different risk for overlapping emission and excitation spectra is when the emission of one fluorophore excites a different fluorophore.

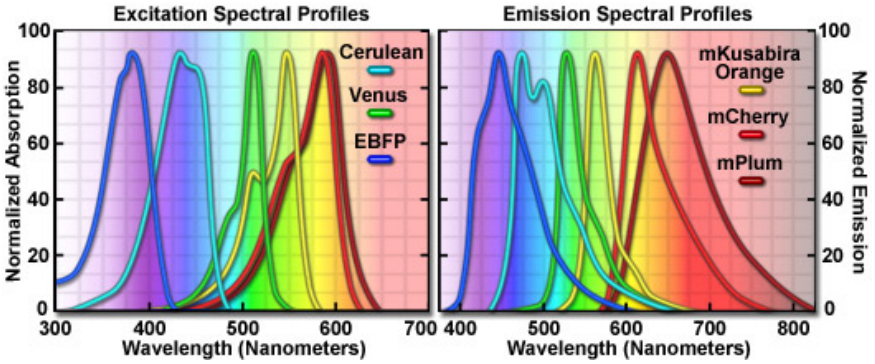


Figure 2.7: Excitation and emission spectra of different fluorescent proteins. Image courtesy of the Olympus microscopy resource center.

2.2.1.3 Confocal microscopy

The 3D contrast of widefield microscopy is significantly improved by the two simple adjustments that are implemented in confocal microscopy. First, the illumination light is focussed by a lens. This causes the light to focus in a region smaller than the resolution of the microscope, thus reducing scattering of the illumination light from nearby points in the sample. Using a moving mirror, this small focussed beam of light is scanned over the sample, such that a full image can be acquired. Fig. 2.8 shows a schematic overview of the illumination in confocal microscopy. To filter the emitted light from the excitation light, a dichromatic mirror is often used. This is a mirror that reflects light at a specific wavelength, e.g. the excitation wavelength, and that passes light at different wavelengths, e.g. the emission wavelengths.

The second important improvement over widefield microscopy is the use of a pinhole or diaphragm in front of the detector. This pinhole partially blocks light originating at different depths than the focal plane. This is illustrated in Fig. 2.9. In this schematic overview, green indicates light coming from a fluorescent protein located in the focal plane, yellow represents light originating under the focal plane, whereas purple corresponds to light from a point above the focal plane. Note that in practice, this light has the same colour, i.e. the fluorophores excite at the same wavelength, but for visualisation we use different colours for the three different situations. Note that light coming from below or above the focal plane is mainly blocked by the pinhole, whereas 100% of the light coming from a point at the focal plain is passed. This allows optical sectioning, i.e. 2D slices of a 3D object can be visualised without having to cut slices of the actual object, thus allowing in vivo 3D imaging.

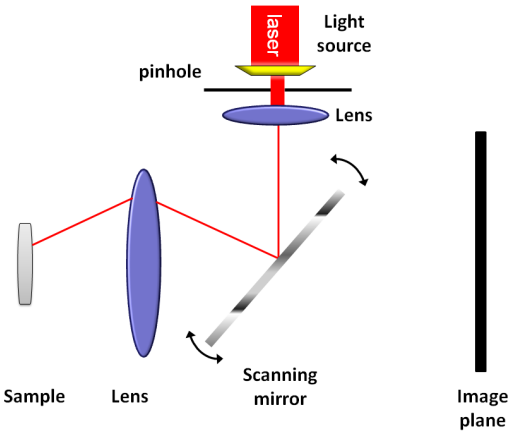


Figure 2.8: A schematic overview of illumination in confocal microscopy.

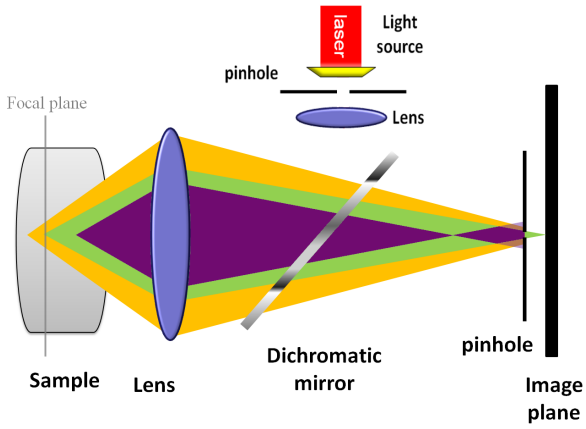


Figure 2.9: A schematic overview of confocal microscopy. Only ligth from the plane of focus can pass to the detector without obstacles. Light from around the focal plain is almost completely blocked by the pinhole in front of the detector.

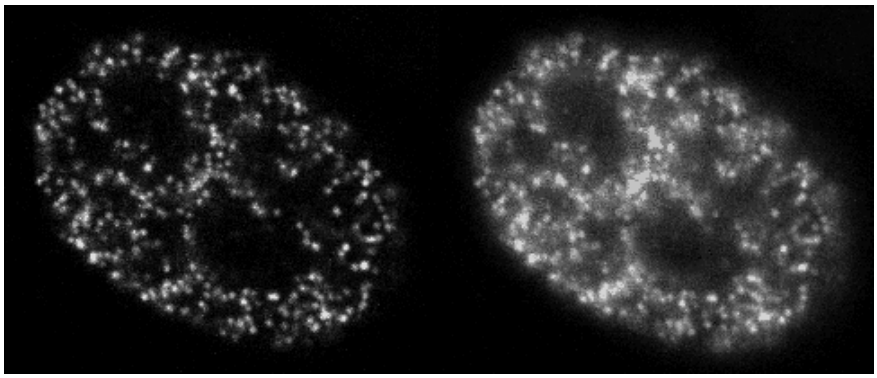


Figure 2.10: Single optical sections of a fixed HeLa-cell nucleus with a scratch replication label highlighting sites of DNA replication. Left: recording in confocal mode. Right: recording in non-confocal mode. Image is courtesy of Steffen Dietzel.

These two adjustments might seem trivial, but they insure that only a single point in the specimen is illuminated and allowed to pass light to the detector. This has a significant impact on the contrast of the image. Fig. 2.10 shows a comparison between confocal and widefield microscopy. This example shows a nucleus with a fluorescent protein marking DNA replication. The left part of the image shows a crisp image acquired with a confocal microscope. The right part of the figure shows the same nuclear section but captured with a widefield microscope. This image is notably more blurred.

2.2.1.4 Optical contrast enhanced imaging

Optical contrast enhanced imaging is a method that exploits differences in the refractive index of different materials to differentiate between structures in the sample. This allows distinguishing between structures of similar transparency without the need of staining. Several microscopy systems use this strategy, but the most popular systems are phase contrast microscopy and differential interference contrast (DIC) microscopy.

Phase contrast microscopy exploits the fact that structures with a different refractive index delay and refract the light by a different amount. The light delay introduced by passing a sample with heterogeneous refractive index results in some waves being phase shifted with respect to each other. Phase contrast microscopy converts these phase shifts in light to brightness changes in the image.

DIC works in a different way, relying on the principle of interferometry to gain contrast. DIC microscopy illuminates the sample by two mutually coherent polarized light beams, which are spatially displaced at the sample plane, and recom-

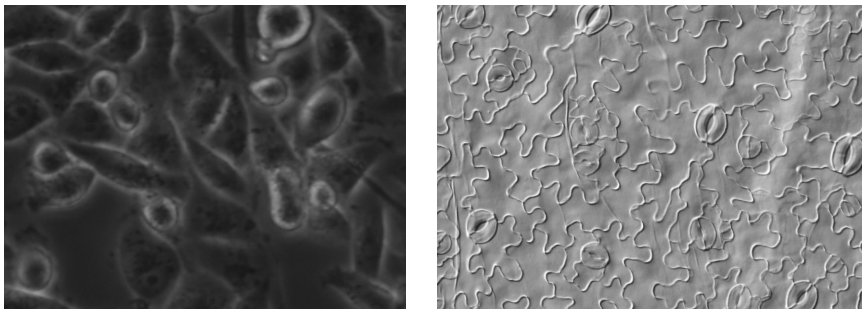


Figure 2.11: Two examples of phase-contrast imaging. Left, an example of phase contrast microscopy of cells. Right, an example of a DIC micrograph of pavement cells of an Arabidopsis leaf.

bined before observation. At recombination, the interference of the two parts, which correspond to two light beams passing in the sample at a small distance from each other, is sensitive to their optical path difference. This is a combination of refractive index and geometric path length. For a more detailed explanation on the optics used in DIC and the analysis of DIC images, we refer to [Davidson 02, Tkaczyk 10, Ochoa Donoso 11].

Figure 2.11 shows an example of both a phase contrast microscopy image and a DIC image. The phase contrast micrograph (left) shows a bright halo around each cell, which is typical for phase contrast microscopy. The DIC image (right) clearly visualises the cell walls, i.e. the double line of white and black. Note that this clear line is mainly visible for cell walls in a certain orientation (the shear direction), i.e. edges going from the left bottom to the right top. Edges orthogonal to this direction, lack contrast. For an explanation of this phenomenon, we refer to [Tkaczyk 10]. The orientation dependent contrast is a small problem for human microscopists, since they can easily connect missing parts of cell walls. However, this presents a serious challenge for automated computer vision methods.

2.2.1.5 Electron microscopy

Instead of using light to illuminate a sample, electron microscopy (EM) illuminates a sample with an electron beam. Since electrons have wavelengths about 100,000 times shorter than visible light photons, EM can visualise much smaller structures then regular optical microscopes, i.e. EM achieves resolution of below 1, whereas optical microscopy is limited to about 187 nm resolution.

Note that electron microscopy is not a single microscopic system, just as optical microscopy is not a single type of microscopes. Instead, it is more a group of

different microscopes, which all make use of an electron beam instead of a photon beam. The types of electron microscopes are:

Transmission electron microscope (TEM): TEM is similar to the bright field microscopic setup (Fig. 2.4), but instead of a light source a beam of electrons is focussed by a set of magnetic lenses onto the sample. This electron beam interacts with the specimen as it passes through. The result of this interaction is then visualised on a fluorescent screen, a photographic film or on a sensor.

Scanning electron microscope (SEM): This type of EM focuses not on what passes through the sample, but what can be measured of the impact on the sample. When the electron beam interacts with the specimen, it loses energy in the form of light, heat, backscattered electrons, etc. By detecting the backscattered electrons, the topography and composition of the samples surface can be modelled.

Serial Block-Face Scanning Electron Microscope (SBF-SEM): SBF-SEM is a technology that provides 3D EM datasets based on a slice and view technique. First, the sample is fixed and embedded in an epoxy resin. From this sample the surface is sliced such that a smooth surface remains that contains no topological information. A stack of images is collected by SEM imaging of the sample followed by slicing with an automated ultramicrotome within the chamber of the SEM. This new technique, proposed in [Denk 04], results in high resolution 3D images, typically under 20 nm even in the z-direction. It is expected that this technology will have a similar impact on biological research confocal microscopy had a few decades ago for the analysis of cells and tissues.

Figure 2.12 shows an example image for each of the previous mentioned electron microscopic systems.

2.2.2 Recent trends and future prospects

2.2.2.1 Increasing spatial resolution

The ability to identify two closely spaced fluorophores is limited by the diffraction of light. The minimal distance, d, between two separate identifiable fluorophores is modelled by Abbe's law:

$$d = \frac{\lambda}{2\,NA}, \tag{2.2.1}$$

with λ the wavelength emitted by the fluorophore and NA representing the numerical aperture of the lens system. The numerical aperture expresses the range of

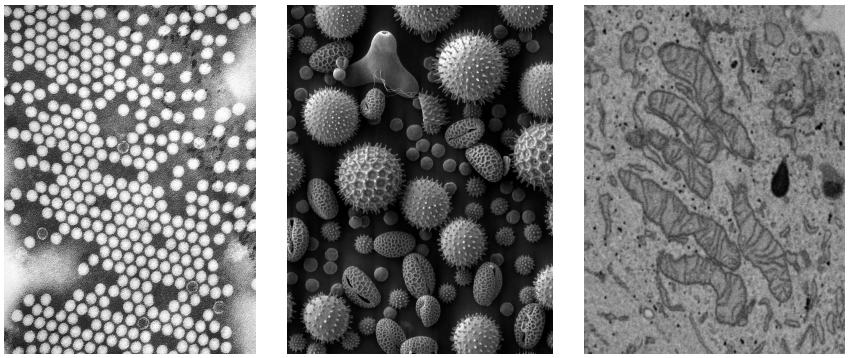


Figure 2.12: Examples of images captured with different types of electron microscopes. Left, transmission electron micrograph of the poliovirus. The image is courtesy of Dr. Fred Murphy, Sylvia Whitfield. In the middle, pollen from a variety of common plants imaged using SEM. Right, mitochondria imaged using SBF-SEM.

angles over which the lens system can accept or emit light. Abbe's law describes the resolution limit of optical systems.

However, eq. (2.2.1) expresses the minimal distance that two emitting fluorophores can be distinguished from each other. But if these fluorophores do not emit at the same time, but after each other, this does not form a problem. In such a situation, the light sensor would detect a bright blob corresponding to the emission of a single fluoreophore at one time frame and a second blob in a next time frame. The blob corresponds to a single light source smaller than the resolution of the microscope, the location of the source can be estimated using special reconstruction methods. These reconstructions are actually probability maps of the location of light sources. A new group of microscopes use this strategy to increase the spatial resolution beyond the diffraction limit observed by Abbe. This technique is called super-resolution microscopy and encompasses PALM, fPALM, STED, SSIM, STORM, etc. One way to achieve to ensure that only a single fluorophore can emit, is to ensure that neighbouring region is unable to emit. This can be achieved by using an extra light beam only on the surrounding region, this approach is used by STED. This extra light will force fluorophores in this region to deplete the excited state before fluorescence takes place.

Fig. 2.13 shows a comparison of an image captured with regular confocal microscopy and with a super-resolution microscope. On the left Vimentin protein are visualised using confocal microscopy, whereas on the right they are captured with super-resolution microscopy. The top part shows the close up of the white square in the central images.

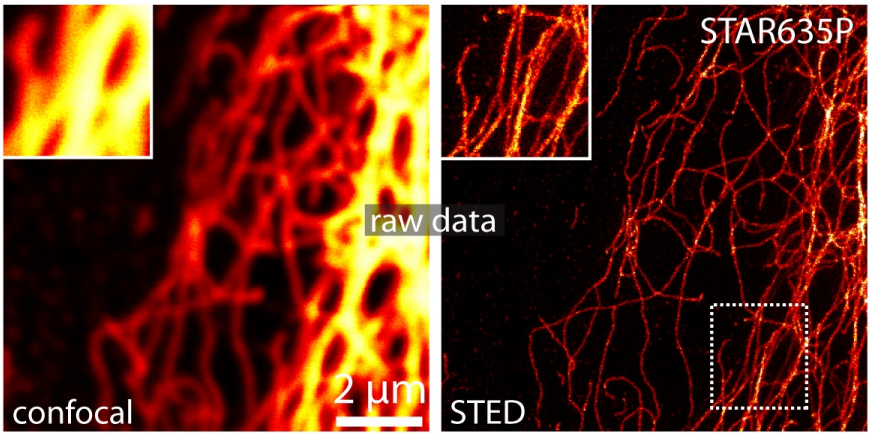


Figure 2.13: Comparison of resolution obtained by confocal laser scanning microscopy (left) and super resolution microscopy (right). The central images each show Vimentin protein, the white boxes are magnified at the top. Image by courtesy of Fabian Göttfert and Christian Wurm, Department of NanoBiophotonics, Max Planck Institute for Biophysical Chemistry.

2.2.2.2 Increasing temporal resolution

Although the use of fluorescent dyes and proteins has great benefit for biological observations, there are still some caveats as well. First, the sample is illuminated using a high intensity laser. This is an unnatural situation for cells, and one can wonder to what extent observations made under such unnatural conditions are actually representative for regular biological behaviour. Limiting the illumination time is good practice for in-vivo studies [Keller 08]. A second challenge is the fact that fluorophores can chemically react on excitation, which can cause the fluorophore to lose its ability to emit photons. This is called photobleaching and this effect is a second reason why fluorescence samples can only be observed for a limited time [Vonesch 06]. Photobleaching and phototoxicity even have been called the Achilles heel of fluorescence microscopy. This has led to several systems that try to reduce the illumination time in order to minimise photobleaching and phototoxicity.

A first approach is to adjust the illumination time based on the sample. A short illumination burst gives an initial estimate of the intensity, and of the presence of a fluorophore at a specific location. Based on the intensity measurement resulting from this initial illumination burst, the excitation-light dose is adjusted to a level such that regions with a low fluorophore concentration still can be observed adequately. Note that for locations with an already high or low intensity little extra illumination is needed, since these indicate locations that either already

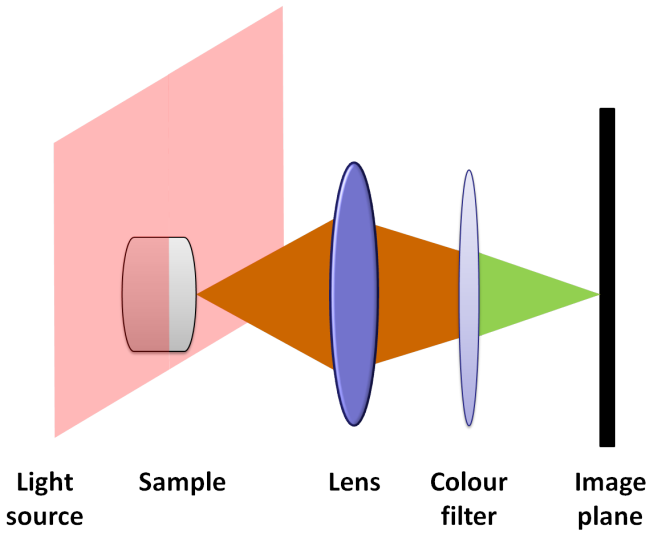


Figure 2.14: A schematic overview of light sheet fluorescence microscopy.

have a signal to noise ratio which is sufficiently high, or to a region where no fluorochromes are present, and as such extra illumination will yield little extra information. This non-linear spatially varying illumination can significantly reduce the overall illumination, i.e. with a factor of 2 to 10 depending on the sample and fluorescent proteins used. This technique is called controlled light exposure time microscopy (CLEM) and already resulted in interesting biological observations [Hoebe 07, De Vos 09a].

A second approach is not to reduce the illumination time, but to use light more efficiently. In most optical microscopic systems the sample is either illuminated from above or below, where if possible the emitted light not originating at the focal plane is blocked by a diaphragm. As such, there is a large part from the sample that is illuminated, without actually contributing to the image, i.e. the part above and below the optical plane. In light sheet fluorescence microscopy (LSFM), the illumination time of a point in the sample is reduced to only illuminating it at the time that it is actually imaged. This is achieved by placing the light source orthogonal to the lens system. Fig. 2.14 shows a schematic overview of LSFM. The full focal plane, and only the focal plane, is illuminated by a sheet of light. Therefore, all fluorophores at the focal plane will emit light that is then focussed by a lens onto an image plane. This approach to reduce light allowed i.a. the study of growing zebrafish embryos, which require 3D 24-hour time-lapse imaging with a temporal resolution of minimal 90s [Keller 08, Keller 10].

A third approach to reduce phototoxicity is called multi-photon excitation. In-

stead of using a single high-energy photon to excite a fluorophore, two low-energy photons are used. Two long-wavelenght photons (typically infrared) must arrive at the same time at the fluorophore, as the summed energy of the two photons will produce the right emission wavelength. This approach is well suited for live cell imaging, since the low energy photons cause less damage to the cells than the short wavelength photons used in confocal fluorescence microscopy.

2.2.2.3 Digital image restoration

Although a lot of research focuses on changing and improving the microscopic setup, improving the images from existing setups is gaining interest from the research community as well. A good understanding of the image degradations that occur in a microscopic setup can help to reconstruct a better image [Rooms 05b, Goossens 10, Aelterman 12, Roels 14]. This typically involves modelling the blur, different types of noise such as Gaussian and Poisson noise, non-uniform lighting, correlation of noise, etc.

Proper digital image restorations can on the one hand enable new analyses that were not possible due to insufficient quality of the original images [Rooms 05a], or can on the other hand be used as a means to reduce the acquisition and illumination time. This would cause the microscope to acquire images of lesser quality, but these degraded images can then be improved using digital image restoration algorithms.

Fig. 2.15 shows three examples of digital image restoration of microscopic images. The left column shows the original RAW images, whereas the right column shows the restored images. The top row shows an example of restoration of a serial block face scanning electron microscopy image. The used restoration technique tackles blur, Gaussian thermal noise, Poisson noise and the correlation of the noise. For more details, we refer to the work of Roels [Roels 14]. The second example shows a 2D maximum intensity projection of a 3D fluorescence microscopic image. The restoration method tackles the prominent blur, which corresponds to a large non-local 3D blur kernel, Poisson noise, thermal Gaussian noise and quantisation noise [Luong 14]. The final example tackles the problem of restoring a fluorescence microscopic image acquired with low illumination and acquisition time, i.e. 5ms compared to 300ms that is commonly used for this data. The restoration algorithm focuses on Gaussian noise removal and intensity enhancement.

2.2.2.4 Multimodal imaging

In macroscopic imaging applications, the use of multiple modalities has already proven its merit, e.g. the combination of IR, chlorophyll fluorescence an VIS im-

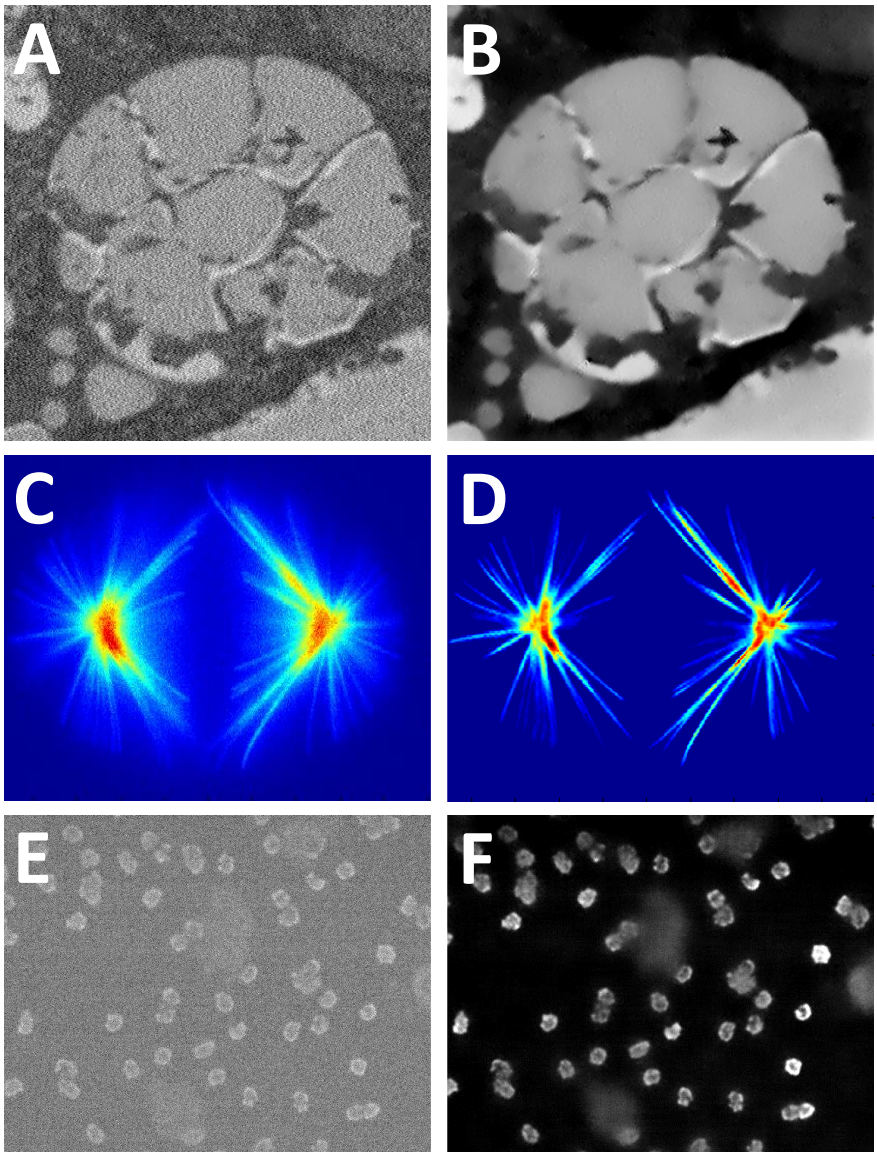


Figure 2.15: Three examples of digital restoration. A) a raw election micrograph, B) the reconstructed image of A [Roels 14], C) a fluorescent microscopic image, D) the reconstructed image of C [Luong 14], E) a fluorescent microscopic image of mononuclear blood cells, F) the reconstructed image of E.

ages. In medical imaging this approach is also gaining interest, e.g. by combining CT, MRI and, or SPECT, a physician gets more relevant information, which allows more accurate diagnosis. Not only more information is available, different image modalities allow better image restoration, i.e. an image of one modality can help to remove noise, or increase the resolution of a different modality [Jovanov 11, Jovanov 13]. There is also a vivid research interest in the use of multimodal microscopic imaging techniques in the domains of forensics [Alsberg 11], plant development [Purbasha 09, Paramita 12], neuroscience [Cote 07, Wilt 09] and pathology [Meyer 11, Miao 12].

This encompasses both inter- and intramodal microscopy. Intra-modal imaging refers to the acquisition and joint processing of multiple image series with the same imaging equipment, thus the same imaging modality is used but with different imaging parameters, e.g. contrast or focus settings [Ochoa Donoso 11]. Inter-modal imaging refers to the combination of different imaging devices, thus applying different imaging modalities. Especially the link between 3D electron and 3D confocal microscopy seems extremely valuable. This would provide both functional information of biological events coming from confocal microscopy in combination with the high resolution enabling the analysis of ultrastructures. Therefore, the next frontier in biological imaging is to develop techniques for Correlative Light and Electron Microscopy (CLEM).

2.3 Conclusion

In this chapter, we discussed the two main groups of biological imaging modalities, i.e. macroscopic and microscopic imaging. For macroscopic imaging in biology the main emphasis lies on high throughput, often for different modalities, i.e. VIS, chlorophyll fluorescence, IR, and recently also hyperspectral imaging. Recent trends include 3D analysis, either by regular 3D camera, e.g. time of flight, or by 3D medical imaging. A different trend is plant analysis in the field, e.g. imaging using drones.

In this chapter, we also presented a number of microscopy systems. This ranges from single lens microscopes; which evolved to systems with multiple lenses and eventually to confocal microscopes; phase contrast imaging and several types of electron microscopes. Each of these systems have their own merit and are therefore still widely in use. We briefly summarised some of the recent trends in microscopy such as increase in temporal and spatial resolution, restoration of degraded images, e.g. due to low illumination and acquisition times, and combined use of different types of microscopes.

This summary of different image modalities is of course not exhaustive but gives

a clear overview of the different types of imaging modalities that are used within biology, and more specific, of the modalities that have been used within this dissertation.

3

Image analysis

Automated methods for biological monitoring have a long history. For example, mechanical devices to measure plant growth, auxanometers, were already in use by the end of the 19th century. Figure 3.1 shows an illustration of such a measuring device. But even back in the 19th century, the visionary German botanist Wilhelm Pfeffer (1845–1920) recognized the potential of imaging techniques over mechanical devices: *"Photographic registration will probably be largely employed in the future, for series of pictures may be obtained which when placed in a kinematograph show the phases of several days' or weeks' growth in a minute or so"* [Pfeffer 03]. Especially the no-contact and non-invasive nature of certain imaging techniques are a major asset for biological monitoring. This has led to the development of high throughput imaging facilities that gather biological relevant data up to the order of dozens of gigabytes a day. Clearly, this requires automated methods to extract biologically relevant measurements from this tsunami of visual data.

Within this chapter, we elaborate on four main groups of image analysis algorithms. We give an overview on the different approaches that can be used to extract relevant information. The methods discussed are chosen based on their importance, their complementarity and or their relevance for the remaining of this dissertation. This chapter is organised as follows. In the next section we explain a rudimentary form of image analysis based on the correlation of simple statistics with certain measurements. In section 3.2 we elaborate on the exploitation of edges to detect

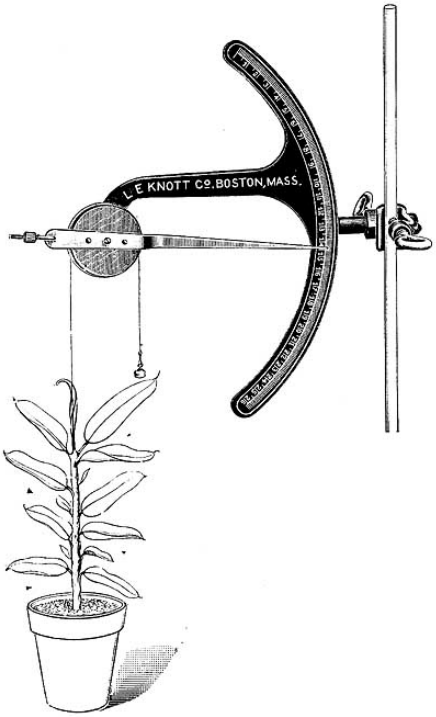


Figure 3.1: An auxanometer is a device for making automated measurements of growth. A figure of an early 20th century auxanometer taken from [Knott 21]

objects of interest. Section 3.3 discusses the use of intensity in order to detect an object of interest. As a final group of image analysis methods we consider model based segmentation, which we discus in section 3.4. As a proof of concept, section 3.5 discusses the experimental results for two different use-cases: the automatic detection of pavement cells and stomata in microscopic images and the growth analysis of full Arabidopsis rosettes. Section 3.6 recapitulates and concludes.

3.1 Correlation based analysis

This group of techniques represents a straightforward approach where simple image statistics, e.g. the intensity histogram or the number of pixels darker than a certain value, are used in a linear regression model to predict a specific biological measurement. This approach has shown interesting results for individual plant phenotyping, e.g. based on top-view and a number of side-view images of a single plant the biomass can be estimated [Berger 12]. It can also be used to calculate



Figure 3.2: A VIS image of fruit bat migrating from deer cave during the early evening, Gunung Mulu national park (Borneo).

group statistics. Figure 3.2 shows an example of such an application. In this image a group of migrating fruit bats are monitored. These bats leave there resting place through a narrow cave exit, as such they are all approximately the same distance from the camera. The number of dark pixels in the image allows us to estimate the number of bats. For example a simple regression model, results in an estimate of 2167 bats in Fig. 3.2. This regression model is trained using a set of small manual annotated images, as a feature the number of dark pixels in the image is used. The benefit of this approach is that it gives an estimate of the number of bats without actually detecting each individual bat itself. There is however a downside to this approach. The regression method fully relies on a training dataset for the calculation of the parameters. If the monitoring conditions differ between training and actual measurements, there is a high risk of erroneous results. The method might fail if for example some tree foliage is captured in the image, but not in the training images, or if the fruit bats leave from a different exit, and as such are at a different distance to the camera, or if the surroundings are darker due to bad weather, etc.

3.2 Edge based segmentation

Edges typically delineate objects in an image, and as such are a valuable asset to detect objects of interest. There are many approaches to detect and exploit edges. In this section we summarise the most important approaches.

3.2.1 Edge detectors

An edge corresponds to an abrupt change in intensity in a certain direction. To automatically detect edges, most methods use the following three consecutive steps:

1. Calculate for each pixel a measurement of change in intensity compared with its neighbouring pixels.

2. Threshold this measurement to make a crisp distinction between potential edge pixels and non-edge pixels.

3. Apply a thinning algorithm to have edges of only a single pixel width.

A straightforward approach for the first step, i.e. to measure intensity changes, is to calculate the magnitude of the gradient of the image. Although different linear filters can be used for the calculation of the gradient, e.g. the Sobel filters, Prewitt filters, Roberts cross, etc., they all have a similar result. The second step of the edge detector, the thresholding has a more significant impact. Especially hysteresis thresholding, which is used in the popular Canny edge detector, shows good results. For more explanation on the different options of thresholding we refer to section 3.3.2. In a final step, the binary image resulting from the threshold step is processed such that the remaining connected components are only one pixel thick. For a discussion on state-of-the-art thinning methods, we refer to [Babin 13a, Babin 14].

Although these edge detectors do not always find closed contours, e.g. the complete boundary of an object, the detected edges can be used as an initial step to detect objects [Kassim 99, Hukkanen 10].

3.2.2 Tracing

Edge detectors mainly rely on the intensity difference at a pixel. These results are often hampered by clutter, which introduce false edges, and by insufficient contrast in combination with noise that often results in parts of edges that are not detected. To improve automatic edge detection, prior knowledge such as the direction or approximate location of an expected edge can be used. Tracing methods are a group of methods intrinsically incorporating different types of prior knowledge [Meijering 10].

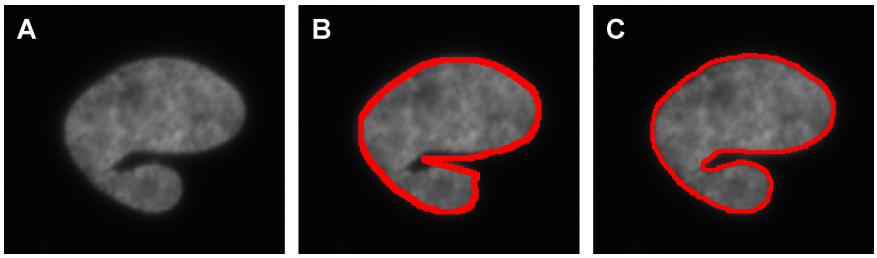


Figure 3.3: An example of contour tracing: A) a RAW fluorescent microscopic image of a pathological cell nucleus, B) an initial estimate of the nuclear contour, C) the contour based on uni-directional line tracing [Cattrysse 13].

3.2.2.1 Line tracing

Uni-directional tracing One of the more powerful prior assumptions that can be exploited is knowledge on the direction of the line that needs to be detected. An example of a situation where the orientation of the edges is known in advance is the detection of a tray in an image. The detection of this tray, can help to locate the actual plants. A different example is where an initial rough segmentation already reveals the overall direction of the edges, but is insufficient to actually detect the lines, e.g. the segmentation of microtubules in a fluorescent micrograph: based on simple thresholding it is possible to detect the orientation, but to accurately detect the microtubules, a different method is needed.

Instead of considering if a pixel belongs to an edge irrespective of its neighbouring pixels, it would be better to combine edge evidence from neighbouring pixels. So instead of only taking the magnitude of the gradient into account, we propose to incorporate information of pixels along the expected edge direction. We will explain this using a simple example where only horizontal lines are detected. We calculate a new edge measurement, $d(.,.)$, starting at the left border of this image. At this border we will, as an initialisation, use the gradient magnitude of the image: $d(i,1) = |\nabla f(i,1)|$, with i the line number and $f(.,.)$ the intensity image. For the second column of pixels we will combine the gradient magnitude and the edge evidence of the already processed neighbouring pixels :

$$d(i,2) = \alpha\,|\nabla f(i,2)| + (1-\alpha)\,\max(d(i-1,1),\, d(i,1),\, d(i+1,1)), \quad (3.2.1)$$

with α a forgetting factor, i.e. a real number less than or equal to one and $d(i,0) = 0$. This can then be repeated for the 3 column, then the 4th column, etc. so in general we get for our new edge measurement

$$d(i,j) = \alpha\,|\nabla f(i,j)| + (1-\alpha)\,\max(d(i-1,j-1),\, d(i,j-1),\, d(i+1,j-1)). \quad (3.2.2)$$

So after this method has iterated over all the columns, the right most column will be a combination of its own gradient magnitude, and edge information coming from all over the image. In general, a column contains information from its own gradient magnitude and from edge information left of the column. As a result, columns at the left, contain little more information then the regular edge detectors that only use pixel per pixel measurements. Therefore we propose to repeat the process, but start scanning from the right column towards the left column:

$$d'(i,j) = \alpha\, d(i,j) + (1-\alpha)\, \max(d(i-1,j+1),\, d(i,j+1),\, d(i+1,j+1)). \tag{3.2.3}$$

This new edge measurement can then be binarised by either thresholding and thinning, conform regular edge detectors, or by calculating a longest path from left to right.

Note that eq. (3.2.2) only combines edge information from its neighbouring pixels, this has a result that only edges with an inclination of less then $45°$ can be detected. So edges that at a certain point have a steeper slope than $45°$ will result in erroneous detections. This drawback is circumvented in [Cattrysse 13], where a variation on the dynamic programming approach allows edges with slope up to $90°$. If the location of the expected edge is approximately known, this constraint can even be further relaxed by resampling the image along an initial estimation of the edge [Mortensen 98, Sun 09]. Fig. 3.3 shows an example of this. The goal is to detect the boundary of the nucleus shown in Fig. 3.3.A. If the approximate location of this boundary is known, e.g. the red curve in Fig. 3.3.B, a neighbourhood around this initial contour can be resampled in such a way that this neighbourhood can be transformed to a rectangle. The discussed tracing method can detect the edge in this rectangle. The detected edge in the rectangle is then projected to the original sample locations in the image. Fig. 3.3.C shows the result of this approach. Therefore, if an initial estimate of the boundary is available, the proposed tracing method can even detect closed contours. For further details on the projection and resampling along a line, we refer to [Kocsis 91, Sun 09].

Multi-directional tracing A line tracing algorithm is useful for applications where a number of isolated lines have to be traced, e.g. sperm tail detection in microscopic images [Bijar 11]. Using standard edge detectors, there is a risk that some parts of the tail are not detected, whereas the previous discussed tracing method would only work if the edge is oriented in a specific direction, which is often not the case. In order to trace a single randomly oriented line, we start from a single point, $\mathbf{s} = (x, y)$ that is known to be part of the edge of interest. This point

is called a seed point. We look for a new seed point in a circular neighbourhood:

$$N_r = \{\mathbf{n} | \mathbf{n} \in \Omega, \text{round}(\|\mathbf{n} - \mathbf{s}\|) = r\} \tag{3.2.4}$$

with $\Omega = [0, K] \times [0, L]$ and $K \times L$ the dimension of the image. The pixel in this neighbourhood that has the strongest magnitude of a gradient, and that does not correspond to a pixel which was a seed point in the past, nor lies in the vicinity of previous seed points, will be chosen as a new seed point, i.e. the next point in the traced edge. This process is iterated until all points in the circular neighbourhood belong to one of the following groups:

- the point, $\mathbf{p}$, has been a seed point;
- the point, $\mathbf{p}$, lies in the vicinity of a previous seed point, i.e. $\nexists \mathbf{v} \in \Omega_s, \|\mathbf{p} - \mathbf{v}\|) < t_1$, with Ω_s the set of all seed points and t_1 a predefined threshold;
- the point, $\mathbf{p}$, has a gradient magnitude less than a predefined threshold, t_2.

The set of all seed points, Ω_s corresponds to the traced line. To get a more densely sampled line representation, interpolation can be used. We illustrate this tracing step using Fig. 3.4.A-C. Fig. 3.4.A shows the gradient magnitude of an image containing cells. Fig. 3.4.B shows a close-up from the red square in 3.4.A. Here the seed point is marked using a green square, and where all the pixels belonging to the circular neighbourhood are delineated in red. This circular neighbourhood is also shown in 3.4.C, where all pixels not belonging to the neighbourhood are set to black. This figure makes it clear why a large circular neighbourhood is necessary. The edges are not a single pixel wide, so if the classical 4- or 8-neighbourhoods would be used there would be a high risk that the algorithm would backtrack some part of the edge, while missing other undetected parts of the edge. For this specific neighbourhood, 4 regions correspond to values above a predefined threshold (marked as a-d). If this is the first iteration, i.e. if no pixel has been a seed in the past, then the white pixel in region c will be chosen as next seed point. However if this is not the first iteration and if a pixel from region c corresponds to a seed point from a previous iteration, then we remove the complete region as candidate seed point. Then the white pixel in region b will have the highest intensity, thus becoming the next seed point.

3.2.2.2 Network tracing

Previous methods have their merit, but they are unsuitable for applications with lines that bifurcate, such as is the case in dendritic branches in neurons or in cell

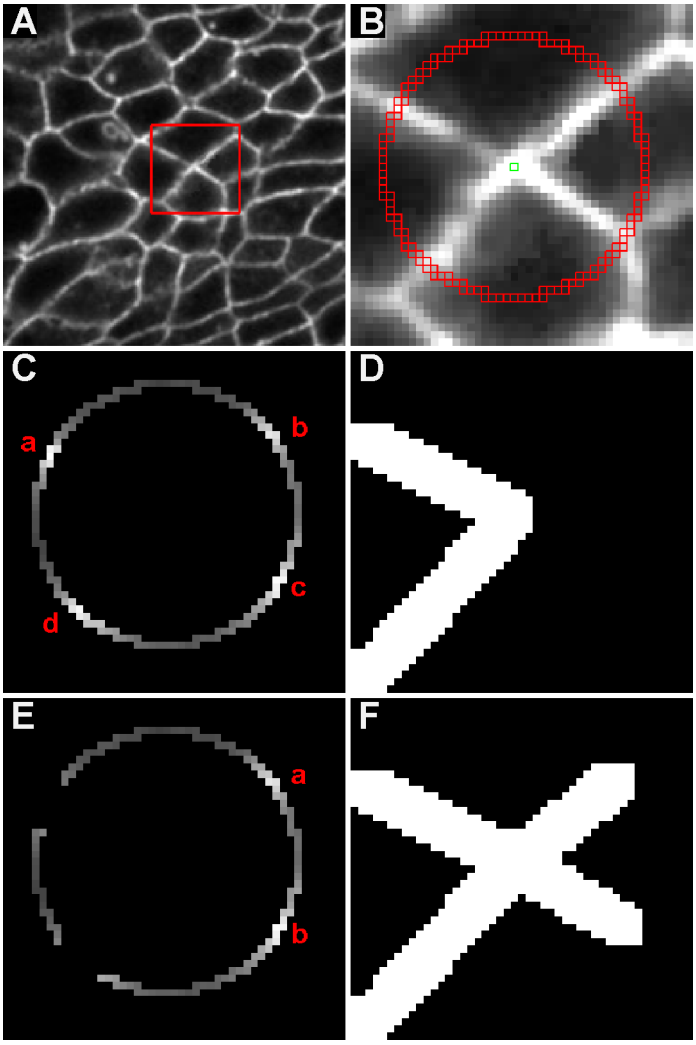


Figure 3.4: An overview of the work-flow of the network tracing algorithm. (A) a microscopic image where cell walls are tagged using fluorescent markers. (B) A close up of the red square in (A). The seed is shown in green and the circular neighbourhood in red. (C) The intensities of the circular neighbourhood. (D) the segment pixels at the start of processing the seed. (E) The circular neighbourhood with the detected pixels removed. (F) the segment at the end of processing the seed [De Vylder 13c].

wall networks (Fig. 3.4.A). In order to extend the tracing algorithm to networks, we start from a segment D_0 that contains a single pixel corresponding to a seed, $\mathbf{s}$, and a circular neighbourhood N_r. In contrast to the line tracing method, the future seed selection procedure will be handled in a different way, instead of only selecting a single new seed, multiple seeds can be chosen. First, the values of all pixels in the neighbourhood, which belong to already detected segment, or are in the vicinity of a segment pixel, are discarded as potential new seeds. Of the remaining pixels in the neighbourhood, the pixel with the highest value is chosen as a potential seed point, $\mathbf{s}_p$. If the value of $\mathbf{s}_p$ is above a threshold, the point is added to the set of future seed points, F. All points on the line between $\mathbf{s}$ and $\mathbf{s}_p$ are added to the segment pixels. In real applications, a line typically is more than one pixel width, thus to avoid backtracking in the tracing algorithm, we also add all the pixels within a small distance, t_d, to the line segment. Therefore, the new set of segment pixels is defined as:

$$\begin{aligned} D_{i+1} \quad = \quad & D_i \cup \{\mathbf{x_2} | \exists \alpha \in [0,1], \exists \mathbf{x_1} : \\ & ((\mathbf{x_1} = \alpha \mathbf{s} + (1-\alpha)\mathbf{s}_p) \wedge (\|\mathbf{x_1} - \mathbf{x_2}\|_2 < t_d))\}, \end{aligned} \tag{3.2.5}$$

with D_i the set of pixels belonging to the already detected segment before adding the new seed. In this formula $\mathbf{x_1}$ corresponds to the points on the line segment between the seed point $\mathbf{s}$ and $\mathbf{s}_p$, whereas $\mathbf{x_2}$ corresponds to the points that are in the neighbourhood of this line segments.

After adding $\mathbf{s}_p$, the neighbourhood is searched for a new potential seed. Note that that this time we exclude all pixels in the circular neighbourhood that also are in D_{i+1}, thus we also exclude pixels in the vicinity of $\mathbf{s}_p$. If the newly found potential seed point has an intensity value higher than a predefined threshold, it is added to the set of future seed points, F. This process iterates until no points in the neighbourhood have a value above the threshold. From this point on, we select a new seed point from the future seed set, F. For this new seed a new circular neighbourhood will be calculated, and new potential seed points will be searched. Seed points for which this processing step does not result in any future seed point are added to the end-point set, E. This iterative seed point selection ends if there are no points remaining in the future seed set F. Note that since points cannot be connected to an already detected pixel, that the end points (the points in E) are only connected to the network as a line. In order to connect the end of the line with the rest of the network, the line tracing algorithm described in subsection 3.2.2.1 can be used.

As an example we continue to work with Fig. 3.4.C. Note that in this neighbourhood, there are 4 regions which correspond to a value higher than the set threshold (denoted with a-d). If Fig. 3.4.D is a visualization of the already detected pixels, than the values corresponding to segment in the neighbourhood are set to zero,

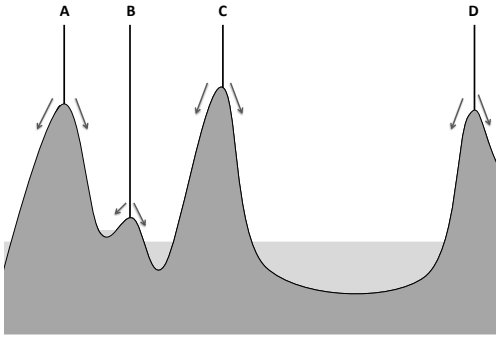


Figure 3.5: A one dimensional example of watershed segmentation, where the actual watershed are marked with a solid black line, the topological minima correspond to light grey areas.

which results in Fig. 3.4.E. After two iterations all possible seed points, a and b, are found, resulting in the segment shown in Fig. 3.4.F. The found seeds, a and b will be processed in future steps.

3.2.3 Watershed

In contrast to previous methods, watershed does not try to detect edges, but instead partitions the image in different segments based on the intensity gradient [Vincent 91]. The watershed algorithm considers the gradient magnitude of the image to be a topological relief, where the gradient magnitude corresponds to altitude in the relief. If a raindrop would fall on this relief, it would flow towards the minimum which lies at the end of the path of steepest descent. All pixel locations from which a drop of water would flow to the same minimum are considered to belong to the same segment, i.e. the segment boundaries coincide with the topological relief watersheds. Fig. 3.5 illustrates watershed segmentation using a one-dimensional example. The topological relief is represented in dark grey, whereas the minima are flooded in light grey. Raindrops falling on this relief will flow towards these minima according to the gradient, i.e. in the direction of the grey arrows. This example relief has four watersheds (a-d).

The watershed is often used as an initial segmentation step that is further refined based using more complex segmentation algorithms [De Bock 06a]. Also for the splitting of a segment into individual objects, the watershed algorithm is a popular choice, see chapter 6 for more details. Aside of an intuitive interpretation, watershed segmentation has the benefit of being computationally and memory efficient [De Bock 05, De Bock 07].

3.3 Intensity based segmentation

Intensity, either in a grey scale image or in a specific colour channel, is another property widely used for image segmentation. This can either be used in a global way, e.g. for the detections of objects with a similar intensity, or based on local statistics, e.g. for the detection of objects that are darker or brighter then their surrounding. In this section we will first start with blob detection methods, which only exploit local intensity features. Next, we describe thresholding methods that can use intensity both in a global or in a local way. Finally, we discuss segmentation methods based on clustering, which use intensity statistics of the full image and which are intrinsically global methods.

3.3.1 Blob detection

Bright or dark blobs are a common appearance in biological imaging. Especially in fluorescence microscopy where fluorescent dyes are used to mark specific proteins or specific biological structures. These fluorescent markers often occur as bright blobs in the images. Telomeres, the ends of microtubules, nucleoli, etc. are just a few examples of structures that appear as bright or dark blobs in fluorescent images. Blob detection is not only popular for fluorescence imaging, but also has many applications on macroscopic imaging, such as fly tracking, virus resistance analysis of plants, the monitoring of mould cultivars, etc.

3.3.1.1 Detection based on linear filters

A simple, yet efficient way of incorporating information from a local neighbourhood is by using linear filters. Some of the most popular blob detectors use a Gaussian filter, i.e. an intensity image, $f(.)$ is convolved with a Gaussian kernel:

$$g(\mathbf{x}, t) = \frac{1}{\sqrt{2\pi}t} e^{-\frac{\|\mathbf{x}\|}{2t^2}}. \tag{3.3.1}$$

This Gaussian filtering can be used to define a scale space representation of the image:

$$L(\mathbf{x}; t) = g(\mathbf{x}, t) * f(\mathbf{x}). \tag{3.3.2}$$

Blobs can be emphasized by calculating a scale-normalised Laplacian on this scale space:

$$\nabla_t^2 L(x, y; t) = t\frac{\partial^2 L(x, y; t)}{\partial x^2} + t\frac{\partial^2 L(x, y; t)}{\partial y^2}, \tag{3.3.3}$$

where bright blobs of approximately the diameter of $\sqrt{2\pi}t$ will result in a high negative value and dark blobs will result in high positive value. To detect blobs of a random size, the following measurement can be thresholded:

$$\mathrm{LoG}(\mathbf{x}) = \arg\max_t |\nabla_t^2 L(\mathbf{x};t)|. \tag{3.3.4}$$

This popular blob detector is called the Laplacian of Gaussians (LoG). Based on scale space theory, the scale-normalised Laplacian can also be approximated using the difference of two Gaussians:

$$\nabla_t^2 L(\mathbf{x};t) \approx \frac{t}{\delta t}\big(L(\mathbf{x};t+\delta t) - L(\mathbf{x};t-\delta t)\big), \tag{3.3.5}$$

the blob detector based on this approximation is the called the Difference of Gaussians (DoG) detector and is used among others in the SIFT feature detector [Lowe 04]. For other blob detectors based on linear filters we refer to [Vonesch 06], or to section 6.2.2.

3.3.1.2 Detection based on morphological operators

Blob detection using linear filters requires a scale space representation for the detection of blobs of random size. Some detectors using morphological operators are not hampered by this requirement. In this section we will first define the morphological operators that are used in this detection technique. Next, the algorithm itself will be explained. The discussed technique has already been proven useful for the segmentation of cornea endothelial cells [Vincent 92], but can also be used for detections in fluorescence microscopic images.

Morphological dilation and erosion Morphological dilation of a grey scale image, $f(.)$, by a structuring set B is defined as

$$(f \oplus B)(\mathbf{x}) = \operatorname*{Sup}_{\mathbf{y}\in B} f(\mathbf{x}-\mathbf{y}). \tag{3.3.6}$$

So this operator replaces a value by the supremum in a small neighbourhood, which is defined by the points in the structuring set B. The structuring element can handle any desired shape, e.g. a line, diamond, circle, etc. Figures 3.6 shows an example of dilation. The left part of Fig. 3.6 shows a grey value image. In this image, a black pixel has value one and a white pixel corresponds to a value zero. This image is dilated using a structuring set that contains all points of a rectangle centred around the origin. The right part Fig. 3.6 shows the result of this dilation.

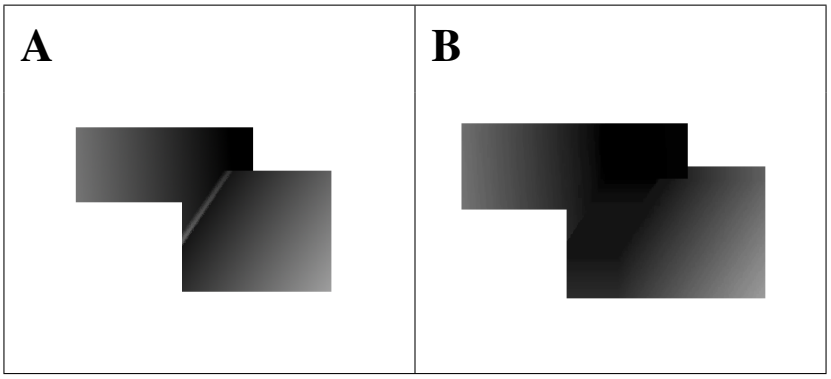


Figure 3.6: An example of morphological dilation. A) a greyscale image which is used as an input Image. B) The dilation of A using a rectangular structuring element.

This results in bright regions being replaced by darker values present in the rectangular neighbourhood defined by the structuring set. This results in an object that is wider then in the original image, also the fine bright structure at the intersection of the two rectangles is removed. The removal of small irrelevant structures is a simple example of the use of morphological dilation, e.g. to digitally remove some soil speckles or veins in a VIS image showing a leaf. To avoid the global growing of the object, another morphological operator could be used consecutively: a morphological erosion. This is similar as dilatation, but instead of the supremum, it uses the infimum. The combination of morphological dilation and erosion is called morphological closing, whereas the subsequent application of erosion and dilation is called morphological opening.

Morphological reconstruction The dilation operator can be used to define other morphological operators such as morphological opening or reconstruction. Morphological reconstruction is useful but little-known method, it can be used to filter out spurious high intensities (e.g. filter away telomeres, without adjusting other boundaries in a fluorescent micrograph), select a small number of objects of interest based on markers, detect local maxima (e.g. bright blobs on a dark background), fill holes (e.g. nucleoli in a fluorescent micrograph), etc. Morphological reconstruction is an operator that processes an image, $f(.)$ based on a structuring set, B and a second image: the mask, $m(.)$. On the input image a morphological dilation is applied. If the pixel value of this dilation is higher than his corresponding mask value, its pixel value is replaced by the value of the mask, i.e.

$$\Big((f,m) \boxplus B\Big)(\mathbf{x}) = \max\Big((f \oplus B)(\mathbf{x}),\, m(\mathbf{x})\Big). \tag{3.3.7}$$

This operator can be repeated n consecutive times:

$$\Big((f,m) \boxplus^n B\Big)(\mathbf{x}) = \max\Big(\big((f,m) \boxplus^{n-1} B\big) \oplus B,\, m(\mathbf{x})\Big), \tag{3.3.8}$$

where $(f,m) \boxplus^1 B = (f,m) \boxplus B$. Note that for a given input and mask image, $f(.)$ and $m(.)$, this morphological operator converges for an increasing n, i.e. $\exists n \in \mathbb{N},\, \forall a \in \mathbb{N}, |\ (f,m) \boxplus^n B = (f,m) \boxplus^{n+a} B$. The result of this convergence is called the morphological reconstruction of an image and is represented as $(f,m) \boxplus^\infty B$.

Figure 3.7 illustrates this concept using a one-dimensional signal. The blue signal is the input signal, and the black signal represents the mask. The red signals correspond to some of the repeated dilations. The green signal shows the actual morphological reconstruction.

H-domes segmentation The morphological reconstruction operator can easily be used for segmentation. In such a case, the mask corresponds to the image that needs to be segmented. From the grey value of every pixel from this image a certain number h is subtracted. This new image is used as an input image for morphological reconstruction. The regional maxima of the image can be extracted by subtracting the reconstructed image from the original one. So for every pixel in the image, the segmentation is defined as follows:

$$\Big((f,m) \boxminus B\Big)(\mathbf{x}) = f(\mathbf{x}) - \Big((f,m) \boxplus^\infty B\Big)(\mathbf{x}), \tag{3.3.9}$$

with all values above zero correspond to segment pixels. eq. (3.3.9) serves as a measure of being a regional maxima: pixels that do not belong to such a maxima are set to zero, whereas the higher the value of this function, the stronger the intensity difference between the pixel and its neighbourhood.

Fig. 3.7 shows the result of this H-domes operator on the right part of the image, where the mask and marker signal (m and f in eq. (3.3.9)) are given. The regional maxima of the image typically correspond to blobs. However, some maxima are the result of noise, or of small clutter such as bleed through fluorescent dye in fluorescence microscopy, etc. To eliminate these segments, an extra thresholding step is applied. Fig. 3.7 shows an example of blob detection using morphological operators, the left part of Fig. 3.7 shows a fluorescent microscopic image where the ends of microtubules are stained using the EB1 fluorescent dye. The right part of Fig. 3.7 shows the actual detections. Note that the blob detection using morphological operators is independent of the scale of objects, i.e. small blobs and large blobs are detected alike, in contrast to linear filter based blob detectors, which typically require the detection of blobs on a Gaussian pyramid.

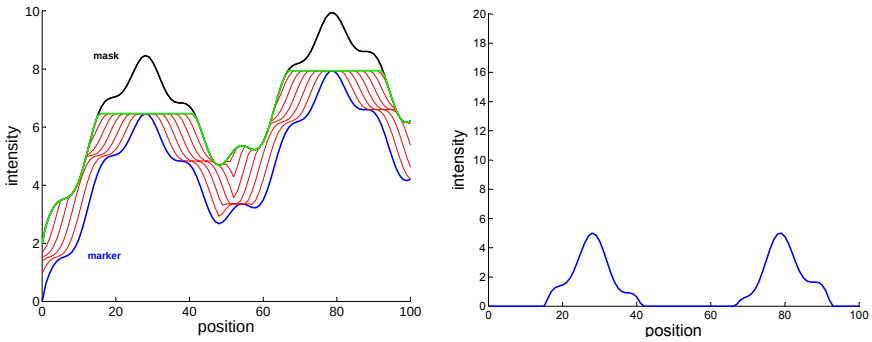


Figure 3.7: An example of morphological reconstruction used for segmentation. On the left, the red curves represent repeated dilatations from the marker constrained by the mask. The green signal corresponds to the morphological reconstruction. The right part of the figure shows the difference of the mask and the reconstructed signal. It is on this signal that a threshold is applied to get the final segmentation .

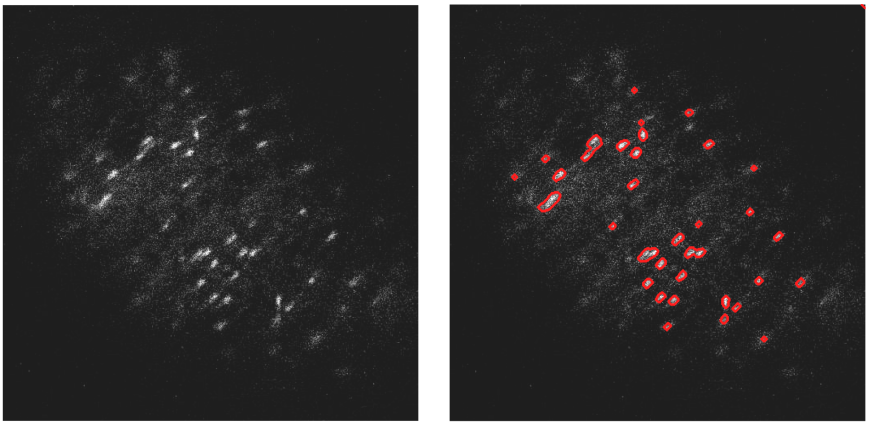


Figure 3.8: An example of the detection of EB1 fluorescent markers using morphological operators [De Vylder 07].

3.3.2 Thresholding methods

Thresholding is one of the most often used segmentation techniques. The main reason for its popularity is that for many applications it results in acceptable results, while still being a very simple, fast and intuitive technique. The concept is simple: pixels with an intensity above a certain threshold, τ are considered to belong to a segment and pixels with an intensity below τ are classified as background, or vice versa. This concept can be further generalised to n different classes:

$$\text{thresh}(f(\mathbf{x})) = \begin{cases} 0 & \text{if } f(\mathbf{x}) < \tau_1 \\ 1 & \text{if } \tau_1 \leq f(\mathbf{x}) < \tau_2 \\ 2 & \text{if } \tau_2 \leq f(\mathbf{x}) < \tau_3 \\ \vdots & \\ n & \text{if } f(\mathbf{x}) \geq \tau_{n+1}, \end{cases} \tag{3.3.10}$$

with $f(.)$ representing the intensity of an image, and 0 to n corresponds to different segments. For many applications the thresholds are set based on some training phase, or are empirically defined [Leister 99, Arvidsson 11]. However this approach is not always suitable for applications where the images can change within the setup, e.g. due to changing lighting conditions, or due to changes in the sample, etc. For such applications, a more dynamic approach is often desired. Estimating proper thresholds based on the images can achieve this. Many heuristics have been proposed in literature to automatically find adequate thresholds [Otsu 79, Abramoff 04, Sezgin 04]. As an example, we describe three of the most popular methods, for a more detailed description on automatic threshold methods we refer to the following works [Otsu 79, Abramoff 04, Sezgin 04]:

Ridler's method: A threshold is iteratively established using the average of the foreground and background class means:

$$\tau^{(t)} = \frac{\mu_f + \mu_b}{2},$$

where μ_f and μ_b are the respective mean intensities of foreground and background based on the threshold from the previous iteration, $\tau^{(t-1)}$. One of the main reasons of the popularity of this method is that it is the default thresholding method in the open source software packet ImageJ [Abramoff 04, Ridler 78]. This software packet is used for a wide range of biological imaging analysis applications.

Otsu's method: This method calculates the threshold by minimizing the variance of the foreground and the variance of the background:

$$\tau = \arg\min_{\tau}(\sigma_f + \sigma_b),$$

where σ_f is the variance of the intensity of all pixels with intensity above τ, whereas σ_b corresponds to the variance of all pixels with an intensity below τ [Otsu 79].

Mixture of Gaussians: The Otsu and Ridler method intrinsically assume that the intensity of both background and object are uni-modal. The main advantage of the mixture of Gaussian approach over the other automated threshold methods, is that it does not make any assumption on the number intensity modes in the image. The peaks or modes in an intensity histogram generally correspond to specific object types in the image, e.g. rosettes will correspond with a peak, whereas soil, trays, labels, etc. typically will result in different peaks. To detect these peaks it is possible to model the histogram, $h(.)$, as a mixture of Gaussians, i.e. to approximate the histogram as a weighted sum of Gaussian probability density functions :

$$h(.) = \sum_{i=1}^{m} \alpha_i N(\mu_i, \sigma_i^2), \tag{3.3.11}$$

where α_i represents a weighting parameter and $N(\mu_i, \sigma_i^2)$ stands for a Gaussian probability density function with mean μ_i and variance σ_i^2. For each type of object an extra Gaussian can be added. The optimal parameters, i.e. weights, means and variances that result in the best approximation of the histogram can be calculated using expectation maximisation optimization [Bilmes 97]. The expectation maximization method is an iterative optimization technique which alternates between two steps: an expectation step, where the expectation of the log-likelihood of eq. (3.3.11) is evaluated using the current estimates for the parameters and a maximization step, where the parameters are optimized in order to maximize the expected log-likelihood found in the first step. For the exact equations used in each step we refer to [Bilmes 97]. One Gaussian curve can be chosen as foreground, whereas all others will correspond to background. The thresholds are then defined at the intensities where the likelihood of the foreground Gaussian curve crosses the likelihood curves of the background curve, i.e. the sum of the remaining Gaussian Curves [Onyango 01, Onyango 03].

While most of the automatic thresholding methods are calculated based on intensity statistics of the full image, it is rather straight forward to adjust this in order to estimate a threshold based on a local level. A first approach is to calculate a threshold for each pixel independently, where an automatic thresholding method of choice can be used that only uses the intensity statistics of a local patch around the pixel. A second approach to adapt the thresholding method to more local intensity statistics is to work in two steps: first apply an initial global segmentation

step, next refine the segments by calculating a threshold for each region around the found segments. Accurate segmentation results have been reported using this approach for the detection of cell nuclei in images where significant intensity variations are observed between individual nuclei [Cattrysse 13].

The concept of getting a rough segmentation result and only in a second step refining the segmentation is an approach, which is also used in hysteresis thresholding. In this approach an initial threshold is used to identify regions that certainly belong to the segments, in a second step a lower threshold is used. Only the segments that are found by the second thresholding step, and that are (partially) found by the first thresholding step, are retained in the segmentation result. Hysteresis thresholding is i.a. used in the golden standard edge detector proposed by Canny [Canny 86].

3.3.3 Clustering based methods

A different approach is to cluster the pixel intensities into a predefined number of clusters. Specially the k-means clustering approach is a popular method for this approach [Keller 08], but more complex clustering methods such as fuzzy k-means clustering are gaining popularity as well [Despotovic 14]. The benefit of clustering methods is that they are easily extended to incorporate features other than intensity as well. This allows the processing of i.a. colour images. Other features such as spatial location of the pixels can also be incorporated, thus resulting in a segmentation method not only based on intensity, but which is also regularised such that pixels in each others vicinity tend to belong to the same segment. The main downside of the use of clustering methods is that they are sensitive to the initialisation of the clusters.

3.4 Model based segmentation

While edges and intensity are the most common features for segmentation, their discriminating power is often insufficient to obtain accurate segmentation results. By incorporating extra knowledge on the objects of interest and on the background, the segmentation accuracy can be significantly improved. This is especially useful when the type of object that needs to be segmented is known in advance. In bio-imaging the setup is often fixed, which can help to model the objects and background appearance. There are generally three different approaches to incorporate prior knowledge into segmentation methods. We will discuss these different approaches and give some specific examples of segmentation methods belonging to these three groups.

3.4.1 Implicit models

The first approach to model the appearance of the background and object of interest is based on extensive training. This approach works in two consecutive steps. First a wide range of features is calculated on a training dataset. This training data should be annotated such that a set of feature values for the ideal segments and a second set of feature values for the background is available. These two sets of feature values are then used to train a machine learning method such as support vector machines or neural networks. Recently the deep learning method has gained a lot of attention due to its ability to implicitly model global features such as shape. In deep learning, first local features are trained using an unsupervised method, i.e. for this training phase no annotated data is needed. The local features detected by this initial step are then fed to new supervised training layers, which train the global features. For a more thorough discussion on implicit models and the use of machine learning techniques, we refer to [Bishop 06, Duda 12].

3.4.2 Constrained methods

In this approach, a model is used more explicitly then with the machine learning methods. For clarity we give some simple examples of what can be considered as explicit models: the area of the segment is constrained to a predefined range, the average segment intensity is above a certain threshold, the decomposition in eigenshapes should fall within two standard deviations of the expected values, etc. For an elaborate discussion on shape models, which are often used as constrained methods, we refer to chapter 8. These explicit models are used to define what a segment should look like. This can be used to analyse and improve an initial segmentation result. Constrained model based segmentation methods can be divided in two different groups: a group that is used more as a post-processing step, and a second group that iteratively updates the segmentation result. The first group checks each segment found by an initial segmentation method. Each segment that does not fit the shape or appearance model is discarded. The benefit of this approach is that few erroneous segmented objects are remaining, which makes it interesting to calculate accurate populations statistics on these segments. Therefore, this is useful for applications where the average appearance or behaviour of an object is more relevant, while an individual object is less of interest [Ochoa Donoso 10, De Vos 10b]. In the second group of methods, each segment is also checked to see if it fits the model. Segments that do not fit the model, are adjusted to a "similar" segment that does. These updated segments not necessarily coincide with an object in the image. Therefore, a new update step is used to adjust the segment to the image. This new segment might not agree with the model, hence this segment is again checked and if necessary, projected to a segment that fits the model. These

steps are repeated until all segments fit the model. Methods belonging to this approach are the active shape model, the active appearance model, super-ellipses, ... [Cootes 95, Cootes 01, Radke 04, Gong 04, Freedman 05, Kauffman 09].

3.4.3 Energy optimisation methods

The constrained based approach applies a black and white strategy: a segment either fits the model or it is does not. All segments that fit the model, according to some criterion, are equally likely, whereas segments that do not fit the model are impossible. To overcome this crisp distinction, a different strategy has been proposed that calculates an energy value for all segments. The energy function usually consists of two parts: a data-fit term and a regularisation term. The data-fit term scores how good or bad the segment delineates specific features in the image, whereas the regularisation term indicates the fit of a specific shape or appearance model. The ideal segment should have a low energy. So most energy based methods algorithms will (iteratively) optimise this energy function in order to segment the image. Note that this energy function generally does not exclude specific segments, such as the constrained based methods do. Instead a very high energy value will be assigned to unlikely segments. Graph-cuts, super-pixel segmentation, region growing, active contours, etc. are just a few of the segmentation methods mentioned in literature that fit within this paradigm [Zhu 96, Felzenszwalb 04, Chan 06, Goldstein 10]. Chapters 4 and 5 will further elaborate on methods that fit within the energy minimisation approach.

3.5 Experimental results

Within this section, we will discuss two specific use cases of generic segmentation techniques. We elaborate on a use case for edge based segmentation, DIC image analysis, and a use case for intensity based segmentation, rosette phenotyping. In the next two chapters numerous use cases for model based segmentation will be discussed. For each use case, we briefly sketch the problem, the added value of an automated analysis, the state-of-the-art of segmentation methods for the specific use case and some experimental results for one of the segmentation methods discussed in this chapter.

3.5.1 Use case 1: cell segmentation in DIC images

Cells are the main building blocks of all organisms. In order to understand functioning, ageing and/or growing of organisms, it is of interest to study the cells

and cell organization of these organisms. This is generally done using optical microscopy. Differential interference contrast (DIC) microscopy is a type of optical microscopy that enhances the contrast in the image based on interferometry. DIC allows for the visualization of transparent cell walls, without requiring any staining. Visual inspection and manual processing of these microscopic images is still common in biological research. The process of annotating data is a tedious, time consuming and error prone task. To overcome this bottleneck in the microscopy pipeline, several automated methods based on computer vision have been investigated. An overview of different cell segmentation and tracking approaches for general optic microscopy can be found in [Shamir 10, Meijering 12].

3.5.1.1 Background

The majority of automated cell segmentation methods start from generic segmentation strategies such as watershed, support vector machines or thresholding [Di Ruberto 02, Marcuzzo 09, Kamentsky 11, Ochoa Donoso 11, Pound 12].

The generic methods generally produce poor results, i.e. with under- or over-segmentation. This is due to clutter, blur, artefacts from underlying cell layers, etc. To overcome these problems, an extra post-processing step is often used to merge and/or split the segments in order to get more accurate segmentation results. This post-processing step is based on the intrinsic shape and appearance of the cells. Most of these methods perform well for images containing a sparse group of cells, but are hampered by dense clusters of cells such as is the case for organs where cells (not necessarily of the same type) are organized in a specific structure. More accurate and robust segmentation results are achieved using fluorescent markers [Quelhas 10a] or motion [Roberts 07, Quelhas 10b], however these approaches are less generic, limiting the applicability for specific applications.

While many approaches in literature attempt to detect the cells directly, it is also possible to focuses on detecting the cells indirectly by detecting the cell boundaries. This duality can help cell segmentation in situations where the cell walls form a network structure: all cells (and cell walls) touch each other and a cell cannot be enclosed by another cell. Fig. 3.9(a) shows an example of such a network structure.

Very few methods dedicated to the analysis of DIC images have been reported [Kam 98, Kagalwala 03, Obara 11]. A few papers focus on transforming the DIC images from an “edge detector like image” towards an “intensity like image”. This however is only useful for images containing few non-overlapping cells. For densely clustered cells this approach will only hamper the segmentation because the boundaries between the cells become less clear. The only method known to the authors that actually tackles dense clusters of cells is discussed

in [Ochoa Donoso 11]. This method first calculates an edge map by fusing edge evidence from different DIC images. Based on this edge map, generic segmentation methods are used to actually detect the cell walls.

3.5.1.2 Segmentation of pavement and guard cells

In order to detect a cell wall network we propose the use of the network tracing algorithm from section 3.2.2.2. We test this approach for the analysis of DIC microscopic images. The DIC image modality raises many challenges for generic state-of-the-art segmentation techniques [Ochoa Donoso 11]. An example of such an image is shown in Fig.3.9(a). To overcome the lack of contrast in specific directions, multiple images (each with contrast in a different direction) are combined using the edge filter proposed in [Ochoa Donoso 11]. The edge map resulting from this detector is then used in the network tracing algorithm. Figure 3.9(b) shows the result of the proposed network segmentation technique. It is clear that the method works well for the detection of pavement cells, i.e. the big puzzle shaped cells (delineated in red). The detection of the stomata (delineated in green) is less accurate. This is because the stomata actually consist of two cells, called guard cells. Sometimes both guard cells are detected, while for one stomata only one guard cell was found. However this is less of an issue, since the stomata have a specific size and shape, allowing to merge the guard cells in to single stomata in a post-processing step.

This segmentation result is quantitatively compared with the method from [Ochoa Donoso 11], which is as far as the authors are aware the only other method developed for the segmentation of *A. thaliana* cells in DIC images. The validation is done in comparison with manual annotated ground truth. The segmentation quality is expressed in terms of correct segmentation (C), over-segmentation (O) and under-segmentation (U): i.e. respectively the number of multiple segments corresponding to a single cell and the amount of multiple cells corresponding to a single segment. The results are shown in Table 3.1. Note that pavement cells are generally better detected than stomata. The over segmentation in the proposed method is only due to the detection of individual guard cells, which is normal since they have the same edge properties as pavement cells. The correct detection of guard cells will make it easier to reconstruct the stomata, since their organization is well known. The method discussed in [Ochoa Donoso 11] does not have this advantage, the over segmentation there is completely random and is not due to the presence of guard cells.

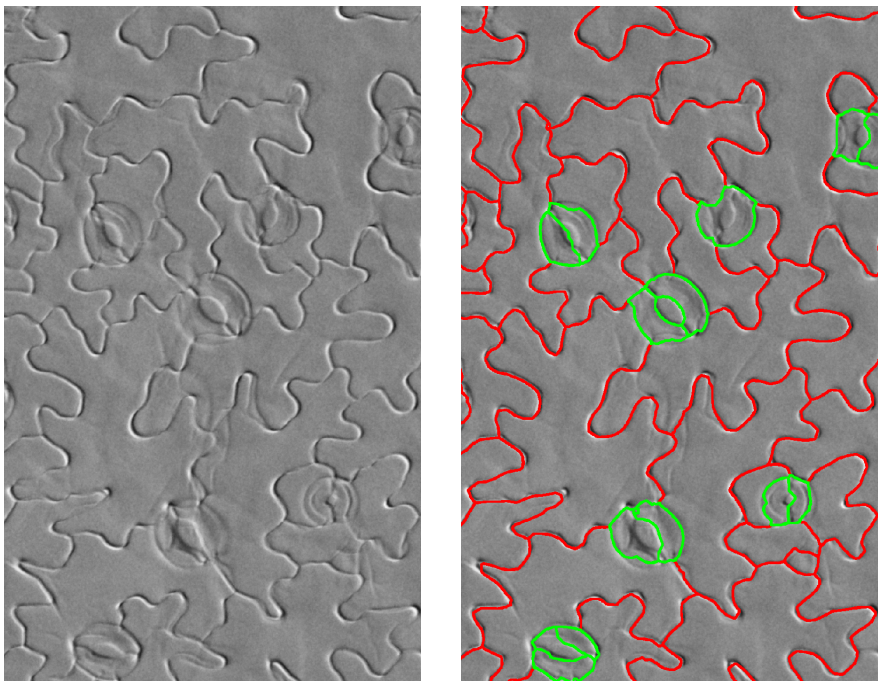


Figure 3.9: An example of cell detection in Arabidopsis leaves imaged using DIC microscopy. Left: the original DIC microscopic image. Right: the result of the proposed method. Pavement cells are delineated in red, stomata in green [De Vylder 13c].

method	cell type	C	O	U
watershed	pavement cells	19	3	1
proposed	pavement cells	21	0	0
watershed	stomata	0	5	3
proposed	stomata	3	0	4

Table 3.1: Quantitative results for the segmentation of DIC images (correct segmentation (C), over-segmentation (O) and under-segmentation (U)).

3.5.2 Use case 2: A. thaliana rosette phenotyping

Integration of tools for simultaneous measurement of plant growth and physiological parameters is a promising way of rapid screening for specific traits. Remote analysis with minimal handling is essential to avoid growth disturbance. For instance, fluorescence imaging can provide information on chlorophyll content without the need for pigment extraction, and at the same time can be used for size estimations of green plants [Jansen 09]. Analysis of plant size or growth is frequently performed by destructive techniques that involve harvesting of whole plants or plant parts at regular time points. This often requires extensive growth room or greenhouse space. In recent years, a number of methods for systematic tracking of plant growth using images have been developed. For growth analysis of the primary root or hypocotyl, Root Trace [French 09] and HypoTrace [Wang 09] were developed primarily for the Arabidopsis community. For analysis of rosette growth, a number of tools have been developed [Leister 99,Jansen 09,Arvidsson 11]. However, since software is tuned for specific hardware, restricting its accessibility for smaller labs. Hardware for large-scale screening is usually based on cameras steered into position with a robot [Jansen 09, Arvidsson 11, Walter 07]. Hence, as many labs do not have the means to invest in such a system, phenotypic analysis of plant growth is still frequently done manually. Plants are either weighed (destructive), scanned (destructive), or photographed followed by manual analysis using image annotation software (non-destructive). These methods are time consuming and call for low-budget, user-friendly alternatives.

3.5.2.1 Background

Several computer vision based methods have been proposed to measure and analyse leaf growth in a non-destructive way. We summarize some of these methods below:

Semi automatic image analysis In [Jaffe 85,Guyer 86,Leister 99] several methods to automatically analyse plant growth are proposed. These methods require the user to either select a set of training pixels or to manually tweak an intensity threshold in order to get robust measurements. The methods are able to detect plant rosettes on a clear background, but are hampered by non-uniform background, e.g. soil [Guyer 86]. This restriction can be solved by illuminating the plants with infrared light.

Motion based methods In [Barron 94, Barron 97, Schmundt 98, Aboelela 05] methods exploiting information from multiple time frames are reported. These

methods show accurate results for high quality images of isolated leafs, but are of course restricted to time lapse sequences. A major drawback of these "optical flow" approaches is that they have difficulties with regions that are partially occluded during a specific time frame.

Vegetation segmentation The following algorithms [Shimizu 95, Onyango 01, Onyango 03, Vigouroux 03] represent a group of algorithms specifically developed for crop segmentation, i.e. automatically delineating a crop from soil in an image. These methods only require a visual (VIS) image of a plant in its natural growing environment. Thus these tools can be used to measure plant growth over time. The vegetation segmentation methods are invariant to light conditions, independent of the camera system and crop size. Unfortunately they rely on the assumption that the images are bimodal, i.e. the images consist of two types of pixels: pixels belonging to a plant or pixels corresponding to soil. The methods will fail if the image contains other types of pixels, e.g. corresponding to a tray, moss, cloth, etc. Although these methods are interesting for specific applications, they are too restrictive for a generic growth analysis system.

Growscreen, growscreen fluoro and lemnagrid The methods in [Walter 07, Walter 09, Arvidsson 11] combine image analysis with specific plant monitoring systems, e.g. the imaging-Pulse Amplitude Modulated (PAM) fluorometer or the Scanalyzer^{HTS} from Lemnatec. All three systems provide a wide range of measurements such as area, relative growth rate, and compactness. With Growscreen Fluoro it is also possible to analyse chlorophyll fluorescence images. The downside of these frameworks is that they only work in combination with specific monitoring systems. A change of tray, camera, focus, lighting conditions, etc. needs a complete resetting of parameters, e.g. thresholds used for the segmentation of rosettes in the work of [Arvidsson 11] are hard coded, i.e. they are not dynamically calculated, instead they have to be manually set, which is tedious and error prone since there are no guidelines on how to set specific parameters.

MRI cell image analyzer [Baecker 07] reports on a general purpose image analysis ImageJ plug in. It is only able to analyse images with a single rosette in a VIS image, which is obviously a great disadvantage when large numbers of plants are to be analysed without an automated/robotised image capturing. MRI Cell Image Analyzer can measure the rosette area over time. All parameters are trained for a specific dataset and should be adjusted for time lapse sequences captured with a different camera system, or with different lighting conditions.

Virtual leaf [Merks 11] proposes a cell-based computer modelling framework for plant tissue morphogenesis, providing a means to analyse the function of developmental genes in the context of the biophysics of growth and patterning. This framework builds a model by alternating between making experimental observations and refining the model. The necessary processing of experiments into quantitative observations is however lacking in this work.

Leaf GUI and limani [Price 11, Dhondt 12b] both provide a framework for extensive analysis of leaf vein and areole measurements. This image analysis framework requires a high resolution, high magnification image of an isolated cleared leaf. Thus, this method is invasive and does not allow to measure the same leaf over time. The framework aims at giving a detailed analysis of the fine structures of a leaf, but does not give any information about the global leaf or plant growth.

3.5.2.2 Quantification of rosette growth using VIS images

We developed an image analysis tool that combines the strong points of the previous methods, while resolving major drawbacks. The tool is based on Gaussian mixture modelling of the histogram, see section 3.3.2. The software tool is able to detect multiple rosettes in an image without assuming bimodal images, i.e. the images can contain parts that are not soil nor plant; nor does it require the use of a specific monitoring system with fixed lighting conditions, trays, resolution, etc. This makes it widely applicable and useful for many different plant phenotyping experiments. The software tool provides a wide range of rosette parameter measurements, i.e. area, compactness, average intensity, maximum diameter, stockiness and relative growth rate (for more details on shape measurements we refer to chapter 8). Apart from analysing regular colour images (VIS images) and/or chlorophyll fluorescence images, the proposed software tool is also able to measure average rosette intensity in thermal infrared images (IR images). To the best of our knowledge, this is the first image analysis tool proposed in the literature offering all these features. Although several image analysis tools have been proposed in the past, a robust measurement tool independent of the monitoring system has not been reported yet. We have developed such a freeware tool to analyse time lapse sequences of *Arabidopsis thaliana* rosettes.

It has been demonstrated that the rosette area is directly proportional to its weight [Leister 99]. Although the relation is not linear for older rosettes as older leaves can become occluded by the newly formed leaves, tracking the rosette area can be a useful tool for monitoring growth. For images in the visual range, current single lens reflex cameras offer a good solution. A 12 Mpixel reflex camera is suitable for the monitoring of 150 (lens 55mm) to 600 rosettes (lens 18mm) in one shot,

from a distance of 2 m. Fig. 3.10.A and B show a picture taken with a handheld reflex camera with a 55 mm lens and the corresponding segmentation picture as generated with our proposed tool respectively. Figure 3.10.C shows the final output as bar graph in a spreadsheet. The ratio between the segmentation result and the manual based measurements are on average 0.97, which is acceptable for the purpose of plant phenotyping based on image analysis.

3.5.2.3 Analysing a composite time lapse sequence of fluorescence and thermal (infrared) images

Fluorescence of chlorophyll can be used to determine the projected leaf area of green rosette plants. Fluorescence images of rosettes are therefore a valuable alternative for VIS images. Using the proposed tool we analysed different rosette plants (Arabidopsis ecotype: Bur-0, Mt-0, Rsch-4, Lip-0 and Ct-1;and ABA mutant rosettes) in time, using a robotised time lapse imaging system containing a fluorescence camera [Chaerle 00, Chaerle 04, Chaerle 07]. It should be noted that similar analysis based on VIS images (e.g. subsection 3.5.2.2) instead of fluorescence images is also possible. The projected leaf area data for ecotypes Bur-0, Mt-0 and Ct-1 confirm the previously observed large rosette size of Bur-0, the intermediate size of Mt-0 and the small size of Ct-1 (Fig. 3.11.A) [Granier 06]. The diurnal difference in area, maximal diameter, stockiness and compactness (for a detailed description on shape description we refer to chapter 8), visible in all lines, reflects up- and downward leaf movements (Fig. 3.11.B,C,D); at mid-day, rosettes are flatter than at midnight [Mullen 06]. Fig. 3.11.E shows the level of small variations in relative growth rate (RGR) when measured with 1h intervals RGR is defined as $\frac{1}{2}ln(\mathrm{Area}_2/\mathrm{Area}_1)$, where Area_2 and Area_1 correspond to the area of the rosette in the current and previous frames, respectively, whereas t represents the time between the two frames. It can be noticed that in some lines, RGR is the lowest at the beginning of the night, in contrast to what was found for circadian hypocotyl extension [Dowson-Day 99]. Analysis of IR images demonstrated that temperatures are the lowest in abi1-1 mutants, indicating strong transpiration (Fig. 3.11.F). The Arabidopsis ecotypes are all warmer than abi1-1. Earlier observations have shown similar transpiration rates for the ecotypes Bur-0, Mt-0 and Ct-0 [Granier 06]. Mutants abi1-1 are known to have reduced stomatal closure and should indeed appear colder in thermal images. By contrast, ABI5 does not influence water loss from plants [Finkelstein 94]; in this image sequence, the values of the 35S::ABI5 are similar to those of the wild type ecotypes. Additionally, diurnal temperature variation is observed, with peaks during the day, due to irradiation heat (Fig. 3.11.F).

These measurements show the usefulness of the proposed tool for plant phenotyping. Furthermore, the measurements are conform to previous observed biological

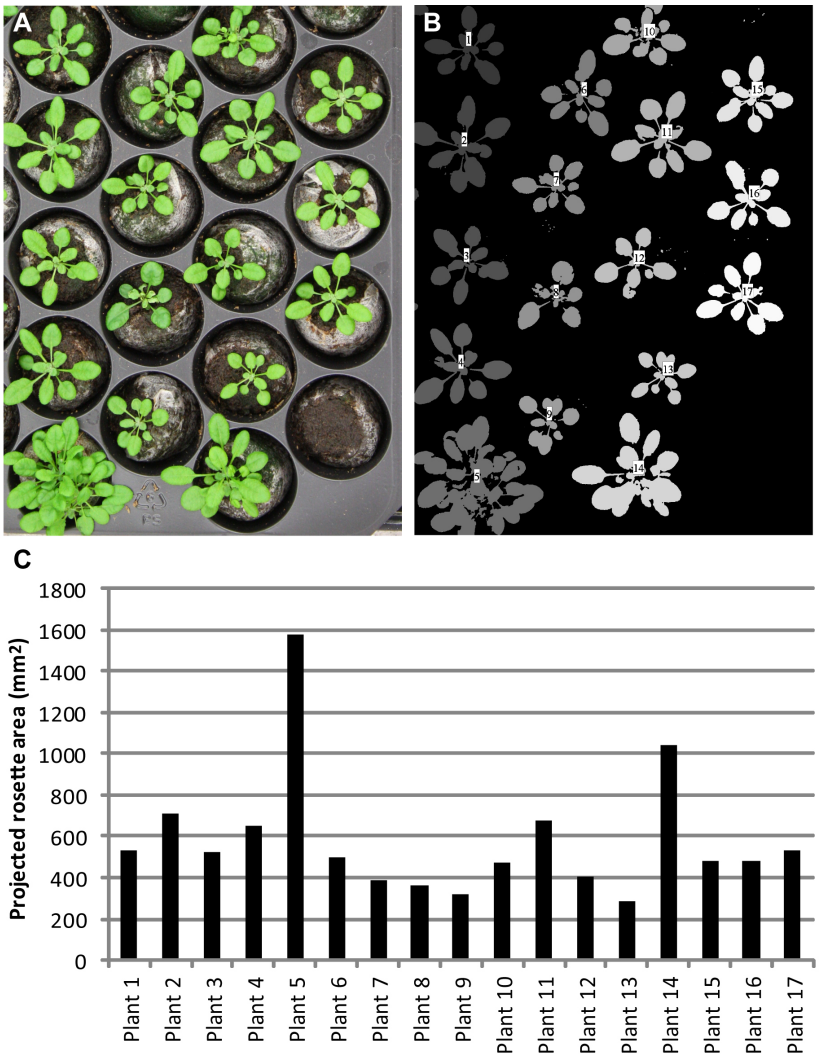


Figure 3.10: Example 1: determination of projected rosette area with Rosette Tracker. A, Original color image showing Columbia-0 plants. B, Segmentation image produced from the color image by Rosette Tracker. Centered numbers in the rosettes from the Rosette Tracker image were repeated on the picture, next to the plant, for clarity. C, Graph showing the projected rosette area per plant from A, as calculated by Rosette Tracker [De Vylder 12c].

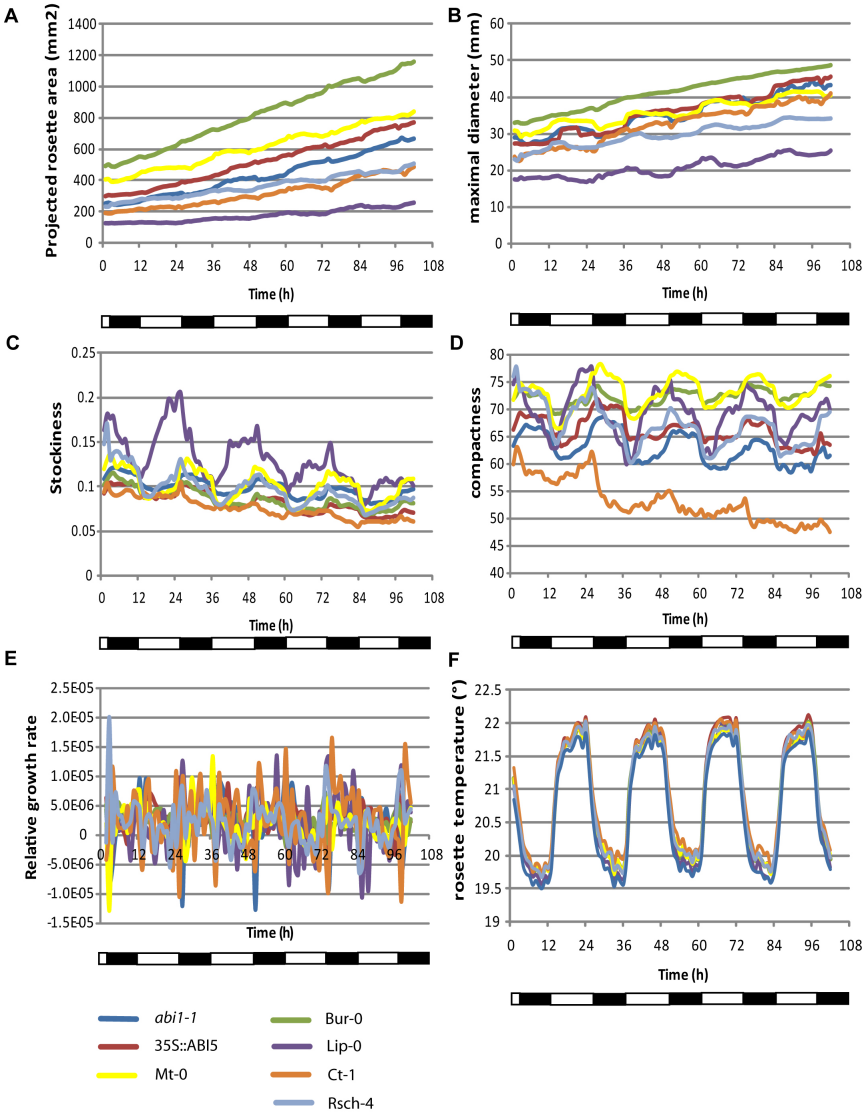


Figure 3.11: Example 2: edited output of results obtained by Rosette Tracker, using a combined analysis of fluorescence and thermal images. 4-week-old rosettes of abi1-1, 35S::ABI5, Bur-0, Ct-1, Mt-0, Lip-0, and Rsch-4 were photographed each hour, and the following average rosette parameters were calculated: projected rosette area (A), maximal diameter (B), stockiness (C), compactness (D), Relative growth rate (E), and temperature (F) [De Vylder 12c].

processes, which indirectly further validates the proposed tool.

3.6 Conclusion

This chapter discussed the problem of image segmentation in biology. It gave an overview of the four major classes of image analysis methods: correlation based methods, edge based segmentation, intensity based segmentation, and model based methods. Furthermore, for edge based segmentation a new tracing method was proposed for the segmentation of network like structures. This method can be used for the detection of cell wall networks in microscopic images, both DIC and fluorescence microscopy.

A freeware toolbox has been developed for the analysis of rosette plants. This tool is not constrained to a single specific monitoring system, can be adapted to different low-budget imaging setups, and requires minimal user input. In contrast with monitoring tools from literature, the proposed toolbox allows to simultaneously quantify plant growth, photosynthesis, and leaf temperature-related parameters through the analysis of visual, chlorophyll fluorescence, and/or thermal infrared time-lapse sequences. The proposed segmentation tool, Rosette Tracker, facilitates the rapid understanding of Arabidopsis genotype effects. The segmentation methods used in this toolbox build upon the methods discussed in the section on intensity based segmentation.

The work presented in this chapter resulted in the publication of a number of international conference papers [De Vylder 07, De Vylder 10d, De Vylder 13c], in an invited talk [De Vylder 12d], and in a peer reviewed journal publication [De Vylder 12c].

Part II

Segmentation & tracking

4

Parametric active contours

The intensity and edge based segmentation techniques discussed in chapter 3 are useful for the analysis of images where the objects are clearly distinguishable, i.e. with clear edges or a substantial difference in intensity between objects and background. Unfortunately, not all images that require automatic analysis have such nice properties. To make a more robust segmentation technique, prior knowledge is typically exploited, e.g. the number of objects in the image. If the segmentation method results in a different number of segments, the result is corrected using splitting and merging operations. Model based approaches typically avoid these complex splitting and merging operations by using stronger prior knowledge [Kass 87, Cohen 93, Cootes 95, Cootes 01, Freedman 05], i.e. by modelling the motion, the appearance or shape of an object. This results in a robust yet general applicable segmentation and tracking framework. A popular group of model based segmentation methods are the active contour methods.

In the active contour framework, a specific energy function is minimised by moving and deforming an initial contour. This energy function should be minimal when the contour is delineating the object of interest. Depending on the application, different energy functions have been proposed. The adaptability of the energy function has resulted in numerous energy functions that can incorporate prior knowledge of motion [Isard 98b, Ray 04, Tang 08], region statistics [Chan 99, Mille 09], expected shapes [Rochery 06, Charmi 08], etc. Two main groups can be distinguished in the active contour framework: one group repres-

enting the active contour explicitly as a parameterised curve and a second group that represents the contour implicitly using higher dimensional functions such as level sets. In the first group, also called snakes, the contour generally converges towards edges or ridges in the image [Kass 87, Tsechpenakis 04, Charmi 08]. The second group typically has an energy function based on region properties, such as intensity [Chan 99, Goldenberg 01]. In this chapter we will focus on the first group, i.e. the parametric active contours. In chapter 5 we will elaborately discuss the second group, i.e. the geometric active contours.

This chapter is organised as follows: first, we will discuss the general active contour framework and the current state-of-the-art. In section 4.2 we discuss a novel external force that, in contrast to state-of-the-art external forces, incorporates both edge and intensity information. Section 4.3 describes how the active contour framework can be adjusted for the analysis of time lapse sequences, i.e. video sequences with low temporal resolution. Next, in section 4.4 we propose the use of prior shape constraints on the data-fit term. By incorporating prior shape knowledge, more accurate segmentation is achieved. In section 4.5 we discuss a new external force which is computationally efficient to calculate, and that results in faster convergence of the active contour. In section 4.6 we propose a probabilistic active contour, which is not optimised by minimising an energy function, instead the likelihood of the active contour is maximised. This avoids the ad hoc setting of weighting parameters, instead the framework relies on the use of likelihood functions. Finally, section 4.7 recapitulates and concludes.

4.1 Background

The original parametric active contour model proposed by Kass et al. [Kass 87], defines the active contour as a parametric curve, i.e. for curve parameter s, $\mathbf{r}(s) = (r_x(s), r_y(s))$, which moves in the spatial domain until the following energy functional reaches a minimum:

$$E_{snake}[\mathbf{r}(.)] = E_{int}[\mathbf{r}(.)] + E_{ext}[\mathbf{r}(.)]. \tag{4.1.1}$$

$E_{int}[\mathbf{r}(.)]$ and $E_{ext}[\mathbf{r}(.)]$ represent respectively the internal and external energy of the contour. The external energy corresponds to a data-fit term, i.e. this term is minimal if the image has strong salient features, e.g. edges, at the location of the active contour. Minimising the external energy alone would often result in erroneous segmentation because each point along the contour is optimised independently. Especially in the presence of clutter and noise the optimisation of individual points can result in suboptimal segmentation. To overcome this, an extra regularisation term is used: the internal energy. The internal energy serves as a

regularisation term that incorporates prior knowledge on the relation between contour points, e.g. neighbouring points in the beginning of the optimisation process should remain in each other's vicinity during the optimisation. As such, contour points are no longer optimised completely independent of each other, which results in accurate segmentation robust against clutter and noise. The active contour framework starts from an initial parametric curve, which could be the result of manual annotation, an initial segmentation step or in the case of video sequences from the segmentation of the previous frame. In the next two subsections, we will elaborate on different internal and external energies.

4.1.1 Regularisation

Many internal energies have been proposed to regularise the optimisation of the active contour. There is however no gold standard: the ideal internal energy depends on the application. When the objects of interest have a specific shape, strong prior knowledge can be exploited. For applications with little knowledge on the object shape, more generic internal energies can be used. In this section we will discuss three types of internal energy, ranging from very generic to more restrictive energy functions.

4.1.1.1 Generic internal energies

One of the most generic internal energy functions is based on the first and second derivatives of the active contour. This internal energy was originally proposed in [Kass 87], and is one of the most popular energy functions due to its simplicity and wide applicability. The influence of the first and second derivative is tuned using two positive weighting parameters, α and β. This results in the following energy function:

$$E_{int}[\mathbf{r}(.)] = \int_0^1 \alpha \, |\frac{d\mathbf{r}(s)}{ds}|^2 \, ds + \int_0^1 \beta \, |\frac{d^2\mathbf{r}(s)}{ds^2}|^2 ds. \tag{4.1.2}$$

The first term, also known as the tension energy, prevents isolated contour point from getting stuck in local optima by penalizing *stretching* the contour. This term corresponds to the length of the active contour. The second term, known as the bending energy, measures the smoothness, e.g. by penalizing sharp corners. This term measures the local curvature of the active contour. Fig. 4.1 shows an example of energy optimisation using different weighting parameters. All active contours are initialised at the leafs boundary. These active contours are solely optimising the internal energy, i.e. there is no data-fit term active. As a result all these contours would eventually converge to either a point or to a circle. Therefore, we ended

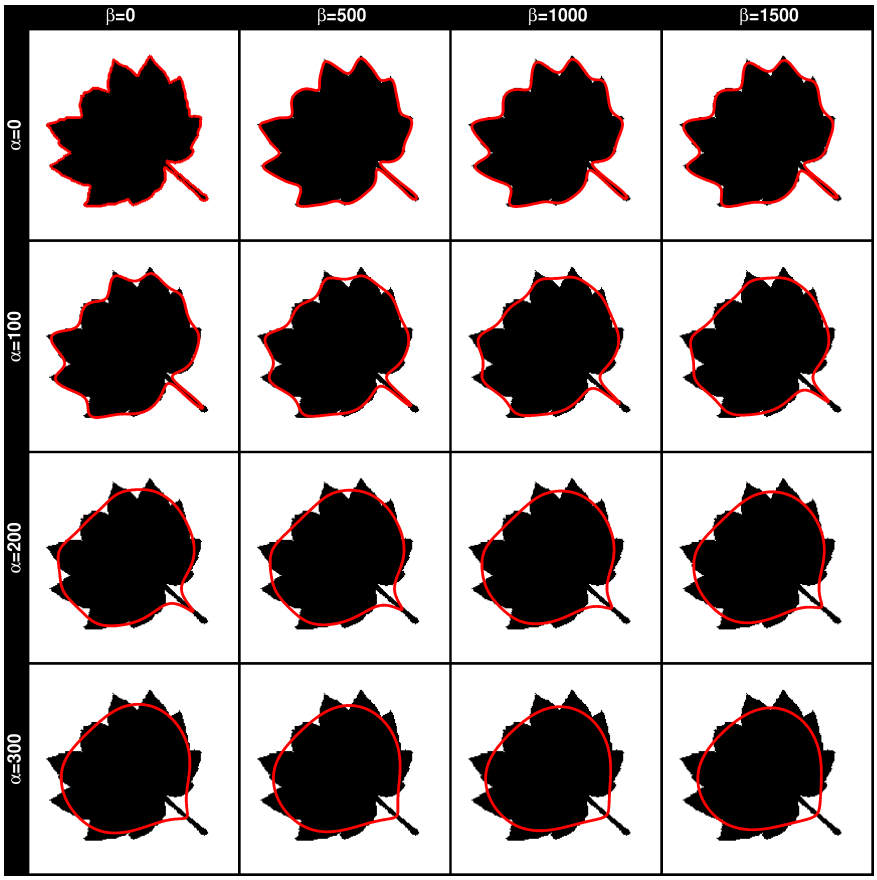


Figure 4.1: An example of active contour internal energy optimisation with different parameter settings.

the optimisation prematurely after 140 iterations. This shows the impact different parameter values have on the active contour. In this image, it is clear that a stronger bending energy, i.e. with a higher β value, results in a smoother contour, with less sharp angles. This energy would converge towards a circle. The tension energy, manipulated using α, minimises the perimeter of the active contour and would eventually converge towards a single point.

Minimizing eq. (4.1.1) for the internal energy defined in eq. (4.1.2), results in solving the following Euler-Lagrange equations:

$$\begin{aligned} \alpha\frac{\partial^2 r_x(s)}{\partial s^2} - \beta\frac{\partial^4 r_x(s)}{\partial s^4} + \frac{\partial E_{ext}\big(\mathbf{r}(s)\big)}{\partial x} &= 0 \\ \alpha\frac{\partial^2 r_y(s)}{\partial s^2} - \beta\frac{\partial^4 r_y(s)}{\partial s^4} + \frac{\partial E_{ext}\big(\mathbf{r}(s)\big)}{\partial y} &= 0. \end{aligned} \tag{4.1.3}$$

These equations can be solved by treating $\mathbf{r}(s)$ as a function of time, i.e.

$$\begin{aligned} \alpha\frac{\partial^2 r_x(s,t)}{\partial s^2} - \beta\frac{\partial^4 r_x(s,t)}{\partial s^4} + \frac{\partial E_{ext}\big(\mathbf{r}(s,t)\big)}{\partial x} &= \frac{\partial r_x(s,t)}{\partial t} \\ \alpha\frac{\partial^2 r_y(s,t)}{\partial s^2} - \beta\frac{\partial^4 r_y(s,t)}{\partial s^4} + \frac{\partial E_{ext}\big(\mathbf{r}(s,t)\big)}{\partial y} &= \frac{\partial r_y(s,t)}{\partial t}. \end{aligned} \tag{4.1.4}$$

the Euler-Lagrange equations described in eq. (4.1.3) are then solved by iteratively taking a small gradient descent step:

$$\begin{aligned} r_x(s,t+\delta t) &= r_x(s,t) - \delta t\frac{\partial r_x(s,t)}{\partial t} \\ r_y(s,t+\delta t) &= r_y(s,t) - \delta t\frac{\partial r_y(s,t)}{\partial t} \end{aligned} \tag{4.1.5}$$

with δt a small step. In practice, the contour is not constant, so after each optimisation step, the contour is resampled [Kass 87, Cohen 93]. The optimisation stops when convergence is reached, i.e. if $r_x(s,t) + \delta t\frac{\partial r_x(s,t)}{\partial t}$ and $r_y(s,t) + \delta t\frac{\partial r_y(s,t)}{\partial t}$ are zero.

4.1.1.2 Higher order internal energies

The generic internal energies regularise the active contour by updating a contour point's location based on the location of its neighbouring points. This is achieved by the use of the first and second spatial derivatives, which in the discrete case means that the two previous and consecutive points on the contour are used to update the location of a point on the active contour. Higher order energies generalise this concept: the location of all contour points within a predefined distance are used to update a contour point. An example of such a higher order energy is proposed in [Rochery 06], where contour points in each other's vicinity should have a similar normal to the curve. This is achieved using the following energy function:

$$E_{int}[\mathbf{r}(.)] = -\iint_{[0,1]^2} \langle \mathbf{t}(s),\, \mathbf{t}(u)\rangle h(\|\mathbf{r}(s) - \mathbf{r}(u)\|)\, ds\, du, \tag{4.1.6}$$

where $\mathbf{t}(.)$ represents the normal to the active contour, and $h(.)$ weighs the interactions between contour points, depending on the Euclidean distance between the contour points. An example of such a weighing function is:

$$h(x) = \begin{cases} 2 & \text{if } x < d - \epsilon \\ 0 & \text{if } x > d + \epsilon \\ 1 - \frac{x-d}{\epsilon} - \frac{1}{\pi}\sin(\pi\frac{x-d}{\epsilon}) & \text{otherwise.} \end{cases} \tag{4.1.7}$$

This function gives a high weight for contour points that are close by, i.e. closer than $d-\epsilon$, and zero weight for pixels further away. The last case ensures that there is a smooth transition from 0 to 2. This smoothness facilitates the optimisation. For more details and other examples of higher order energy functions, we refer to [Rochery 06, Aymen El Ghoul 09, Horvath 09].

4.1.1.3 Shape based internal energies

Both generic and higher order energies regularise the active contour based on local features, i.e. in order to correct a specific contour point, only a set of contour points in the vicinity are used. Another approach is to use global features, i.e. use all contour points. By exploiting global shape, more robust and accurate segmentation is possible, at least under the assumption that the global shape of the objects of interest can accurately be modelled. Several internal energy functions that model the global shape have been investigated in literature [Isard 98a, Goobic 01, Derrode 06, Charmi 08]. These generally use a set of shape descriptors, such as Fourier descriptors or eigenshapes. The internal energy is commonly defined as the Euclidean distance between the shape descriptors of the active contour and the shape descriptors of the expected shape. For a detailed discussion on shape descriptors, we refer to chapter 8.

4.1.2 Data-fit

Where the internal energy serves as a regularisation term, the external energy takes the role of a data-fit term. The external energy is derived from the image such that the contour with a minimal external energy should ideally delineate the object of interest. Given a feature map $f(x, y)$, i.e. a map with high values at locations with high evidence for a feature of interest, the external energy is defined as:

$$E_{ext}[\mathbf{r}(.)] = -\int_0^1 f(\mathbf{r}(s))ds. \tag{4.1.8}$$

The external energy defined in eq. (4.1.8) is generic and can use any type of features, e.g. ridges, blobs, corners, etc. One of the most popular used features are edges, which can be modelled using one of the following feature maps:

$$f(x, y) = |\nabla g(x, y)|^2 \tag{4.1.9a}$$

or

$$f(x,y) = |\nabla\big(l_\sigma * g\big)(x,y)|^2, \tag{4.1.9b}$$

where $g(x,y)$ corresponds to a grey level image, and $l_\sigma(x,y)$ represents a 2D low-pass filter, e.g. a Gaussian kernel with standard deviation σ. Note that the active contour will not be attracted towards edges if they are far away.

4.1.2.1 Data-fit with robustness to local optima

Active contours using the external energy term in eq. (4.1.8) usually require a good initialisation, i.e. close to the object boundary, in order to achieve correct convergence. This limitation is caused by the nature of the external energy term, which is hampered by noise, local optima and which is typically non-flat only in the proximity of the object boundary. So outside the vicinity of a boundary, the image gradient has a magnitude of either zero, due to the flatness of these external energy terms, or points in random directions due to noise. This seriously hampers gradient based optimisation of the energy terms. The maximum distance from which a free particle is attracted to an edge is called the capturing range. To extend the capturing range of the external energy defined in eq. (4.1.8), Xu and Prince [Xu 98] proposed the use of an external force field, $\mathbf{v}(x,y)$. This force replaces the partial derivatives of $E_{ext}(\mathbf{r}(s,t))$ in eq. (4.1.4), i.e.

$$\frac{\partial \mathbf{r}(s,t)}{\partial t} = \alpha\frac{\partial^2 \mathbf{r}(s,t)}{\partial s^2} - \beta\frac{\partial^4 \mathbf{r}(s,t)}{\partial s^4} + \mathbf{v}(\mathbf{r}(s,t)). \tag{4.1.10}$$

Replacing the partial derivatives with a more general force field allows optimisation that is less hampered by local optima. Different forces have been proposed for the data-fit optimisation of active contours. We briefly discuss the most popular ones:

Gradient vector flow A useful approach is to calculate a new force field, $\mathbf{v}(.,.) = (u(.,.),\, w(.,.))$, which resembles the partial derivatives of $E_{ext}(\mathbf{r}(s,t))$ in eq. (4.1.4) at locations with high gradient magnitude, and which is smooth in other regions. Such a force approximates the gradient of the feature map near edges, whereas the smoothness helps to overcome the lack of direction at locations where there is no high gradient magnitude. A force field with these properties can be calculated by minimizing the following energy function:

$$E_{GVF}[\mathbf{v}(.,.)] = \iint_{[0,M]\times[0,N]} |\nabla f(x,y)|^2 |\mathbf{v}(x,y) - \nabla f(x,y)|^2 + \mu\Big(\frac{\partial u(x,y)}{\partial x}^2 + \frac{\partial u(x,y)}{\partial y}^2 + \frac{\partial w(x,y)}{\partial x}^2 + \frac{\partial w(x,y)}{\partial y}^2\Big)\, dx\, dy, \tag{4.1.11}$$

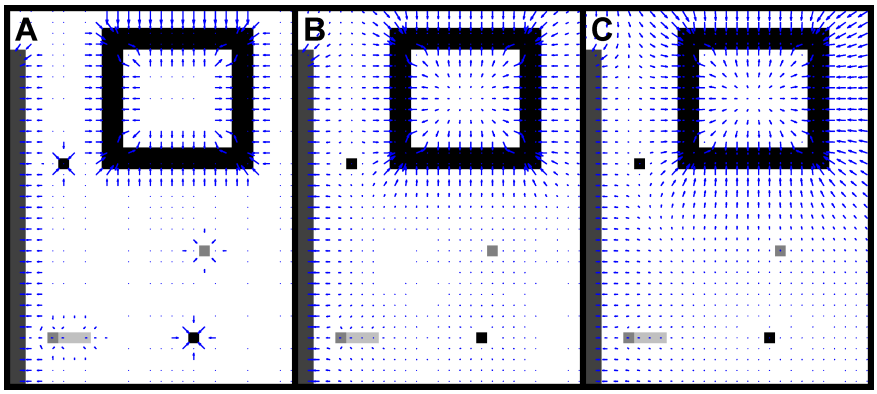


Figure 4.2: An example of a GVF force field using different weights for the smoothness terms.

where $M \times N$ represents the dimension of the image, and μ is a non-negative parameter expressing the degree of smoothness imposed on the field $\mathbf{v}(.,.)$. $f(.,.)$ represents a feature map such as defined in eq. (4.1.9).

The first term of eq. (4.1.11) is a fidelity term, forcing the field $\mathbf{v}$ to resemble the original edge force in the neighbourhood of edges, whereas the second term keeps the field $\mathbf{v}$ smooth. This external force is called *Gradient Vector Flow* (GVF) field. The force field with minimal energy can be found using gradient descent [Xu 98].

Figure 4.2 shows an example of the GVF force field calculated using different weights for the smoothness factor. Fig. 4.2.A corresponds to a force field with a small μ, i.e. with practically no smoothness constraints. This force is very similar to the gradient of the image, which results in a clear force near edges and isolated noise pixels and a force of zero magnitude in uniform regions. Fig. 4.2.B shows a GVF force field with a higher smoothness weight. The smoothness constraint results in force pointing towards edges, even at locations not neighbouring an edge or isolated point. The force magnitude around isolated pixels significantly decreased compared to Fig. 4.2.A. Fig. 4.2.C shows the force field with high weight for the smoothness term. This results in a force field where for each pixel location, the force points towards an edge. Isolated noisy points no longer show any influence at the force field.

Vector field convolution A different approach is proposed by Li et al. [Li 07a] and by Wang et al. [Wang 08], who define an external force by low pass filtering the partial derivatives of $E_{ext}(\mathbf{r}(s,t))$ in eq. (4.1.4). This force is called *Vector*

Field Convolution (VFC), and is defined as follows:

$$\mathbf{v}(x,y) = \Big(\frac{(k * \partial f)(x,y)}{\partial x}, \frac{(k * \partial f)(x,y)}{\partial y}\Big), \tag{4.1.12}$$

where k represents a kernel function. The used kernel function influences the capturing range of the active contour. While technically all kernel functions can be used, kernel functions with values inversely proportional to the distance to the centre seem to be favoured [Li 07a, Wang 08].

Other force fields proposed in literature comprise virtual electric field (VEC) [Park 02, Han 07], charged particle model (CPM) [Jalba 04], boundary vector field (BVF) [Sum 07], etc. Although many different forces have been proposed over the years, it is still the GVF field that is considered to be the gold standard. One of the reasons is the fact that it is simple to tweak the force to a specific application, since it only requires a single intuitive parameter.

4.2 Intensity based data-fit

The use of an external force increases the capturing range of an object and can consequently relax the initialisation constraints. However, when an image contains multiple objects, the external force field can bias the contour evolution towards objects with higher contrast. In Fig. 4.3 we illustrate this problem with an example containing a high contrast blurred square and a low contrast circle. The initialisation of the active contour is shown red. Figure 4.3.B shows the magnitude of the gradient. Note that the square displays a gradient with higher magnitude than the circle, this is due to the better contrast between the black foreground and white background. Figure 4.3.C shows the GVF force field. In the region located between the two objects, the resulting GVF force is mainly directed towards the black square, this is due to the difference in magnitude of the gradient. If the active contour is initialised at a small distance from the circle, the segmentation can yield unpredictable results such as in Fig. 4.3.D.

A simple strategy to solve this problem is increasing the weighting parameters of the internal energy, e.g. α and β in eq. (4.1.2). This will reduce the risk of converging to a wrong optimum under the influence of neighbouring objects. The negative aspect of giving more weight to the internal energy is that it also reduces the possible shapes that can be segmented.

Another possible solution is limiting the influence of objects far away, e.g. by changing μ in eq. (4.1.11) or by choosing a different kernel function in eq. (4.1.12). This strategy however limits the capturing range of an object, thus requiring a good initialisation. This is not always feasible in tracking applications where the

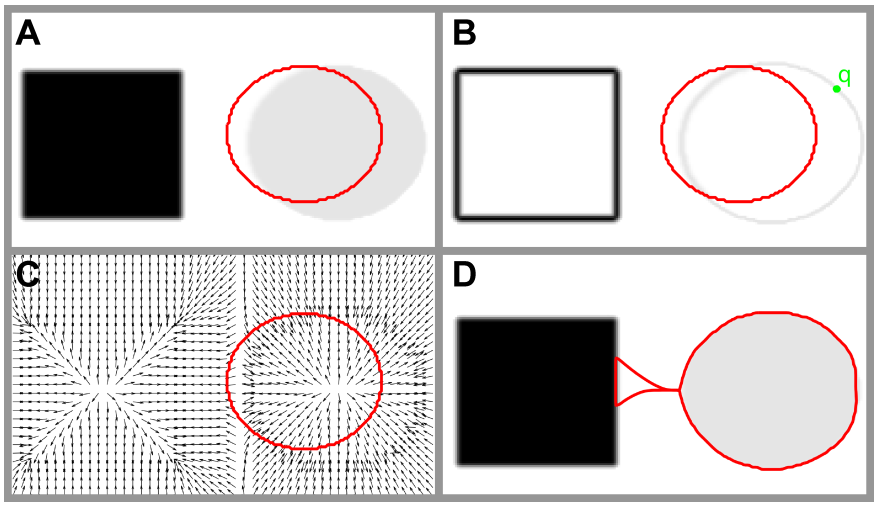


Figure 4.3: An example of erroneous segmentation using a GVF active contour. A) the input image, containing two objects and the initialisation of the active contour shown in red. B) the magnitude of the gradient of the image. C) An illustration of the GVF force field, D) the actual segmentation result using the force field of C [De Vylder 09a].

resulting active contour of frame t is used as an initialisation of the active contour of frame $t+1$. So the GVF force should range as far as the objects of interest can move between two subsequent frames. A disadvantage of tracking using active contours is the risk of error propagation: erroneous active contour segments will be used as initialisation in the next frame.

To cope with this problem, literature reports on different manipulations of the external force field such that gradients towards a certain direction are favoured [Cohen 93, Tang 08]. This approach is useful if a good estimation of the object's location is available, since the force field can then be adjusted such that gradients in the direction of the object are favoured. Unfortunately estimating the object's location is not always straightforward, and strong deviations between the estimated location and the actual location will result in erroneous segmentation.

4.2.1 A joint edge and intensity based force field

Instead of manipulating the force field based on an estimated location, we propose to exploit prior knowledge about the intensity of the object. We propose a new external force field that can be used to evolve an active contour towards objects with a predefined intensity. A major challenge is to estimate the object intensity. To make this more concrete, consider the green pixel, q, in Fig. 4.3.B. At this pixel,

there is clear edge evidence, but estimating the intensity of the object it borders is less straightforward. This cannot be estimated at the pixel's location, which has an intensity somewhere between the object intensity and the background intensity. Therefore, we define the neighbouring intensity operator. To emphasise that this is an operator, and not a function we will use the $\diamondsuit_\eta$ symbol instead of a letter. The neighbouring intensity operator is defined as follows:

$$\diamondsuit_\eta g(x,y) = g\Big(x + \eta\frac{\partial g(x,y)}{\partial x}, y + \eta\frac{\partial g(x,y)}{\partial y}\Big). \tag{4.2.1}$$

This operator samples the image intensity, $g(.,.)$, located at a distance $|\eta|$ from the pixel position (x,y) in the direction of the gradient. The sign of η indicates if the sample is taken in the direction of the gradient or opposite this direction. This neighbouring intensity will help us to distinguish between objects of interest, i.e. with the expected average intensity, and clutter. However, before being able to distinguish between edges bordering relevant objects and irrelevant objects, we first need to model an object of interest. We model the object intensity by its likelihood, as an example we propose the use of a normal distribution. So the likelihood that the intensity of a pixel (x,y), under the hypothesis, that the pixel corresponds to an object of interest is modelled as $\mathrm{p}_{\mathrm{G|H}}(g(x,y)) \sim N(\mu,\sigma^2)$.

Using the likelihood of the intensity of an object, combined with the neighbouring intensity operator allows us to calculate the likelihood that an edge pixel borders an object of interest, i.e. an object with the right intensity. Remember that the neighbouring intensity operator allows us to sample the intensity, either in the direction of the gradient, or in the opposite direction. This is because we consider the possibility that the object's surroundings can be darker or lighter than the object itself. It is however sufficient that there is an object in either one of these directions. This results in two hypotheses: a hypothesis that an object of interest is present in the direction of the gradient, H_0, and a second hypothesis that there is an object of interest present in the direction opposing the gradient, H_1. Note that it is irrelevant for an edge pixel in which direction the object lies, i.e. it is sufficient that one of these hypothesis is true, it does not matter which one. This results in the following likelihood:

$$\begin{aligned} \mathcal{L}_i(x,y) &= \mathrm{p}(H_0 \vee H_1 \mid \diamondsuit_\eta g(x,y),\, \diamondsuit_{-\eta} g(x,y)) \\ &= \mathrm{p}(H_0 \mid \diamondsuit_\eta g(x,y)) + \mathrm{p}(H_1 \mid \diamondsuit_{-\eta} g(x,y)) \\ &\quad - \mathrm{p}(H_0,\, H_1 \mid \diamondsuit_\eta g(x,y),\, \diamondsuit_{-\eta} g(x,y)). \end{aligned} \tag{4.2.2}$$

This can be simplified into the following equation using Bayes' theorem, the assumption that both hypotheses are equally likely, and the fact that $\mathrm{p}_{\mathrm{G|H}}(\diamondsuit_\eta g(x,y)) =$

$\mathrm{p}(\diamondsuit_\eta g(x,y) \mid H_0)$:

$$\begin{aligned}\mathcal{L}_i(x,y) &= h\frac{\mathrm{p}_{\mathrm{G|H}}(\diamondsuit_\eta g(x,y))}{\mathrm{p}_{\mathrm{G}}(\diamondsuit_\eta g(x,y))} + h\frac{\mathrm{p}_{\mathrm{G|H}}(\diamondsuit_{-\eta} g(x,y))}{\mathrm{p}_{\mathrm{G}}(\diamondsuit_{-\eta} g(x,y))} \\ &\quad -h^2\frac{\mathrm{p}_{\mathrm{G^2|H}}(\diamondsuit_\eta g(x,y),\, \diamondsuit_{-\eta} g(x,y))}{\mathrm{p}_{\mathrm{G^2}}(\diamondsuit_\eta g(x,y),\, \diamondsuit_{-\eta} g(x,y))}. \end{aligned} \tag{4.2.3}$$

The likelihood that there is an object neighbouring the pixel in a specific direction, i.e. that H_0 or H_1 is true, is represented by a constant factor h. This factor can be estimated using the ratio of the perimeter of the initialisation and the size of the image. The parameter η is set such that it ensures that the neighbouring intensity operator samples pixels of the objects interior, instead of a blurred edge. This parameter can either be fixed, e.g. at a distance of two pixels, or it can be defined dynamically based on the gradient's magnitude.

The likelihood that a pixel corresponds to an edge bordering an object of interest can then be modelled by combining the likelihood that a pixel corresponds to an edge, and that the pixel is at the border of an object of interest. For simplicity we assume that both likelihoods are independent of each other, i.e.

$$\mathcal{L}_{ie}(x,y) = \mathcal{L}_e(x,y)\mathcal{L}_i(x,y). \tag{4.2.4}$$

Within our experiments, we model the likelihood that a pixel is part of an edge, $\mathcal{L}_e(x,y)$ using a normal distribution, although other distributions could be used as well, e.g. in the presence of shadows a more skewed distribution could be better suited. By using this likelihood as a feature map in a force field, i.e. $f(x,y) = \mathcal{L}_{ie}(x,y)$ in eq. (4.1.11) or eq. (4.1.12), we get an external force that evolves an active contour towards edges that are near an object with the right intensity. The proposed likelihood in combination with the GVF force field is called intensity gradient vector flow (IGVF) field.

The proposed feature map, and by extension the proposed force field, assumes prior knowledge of the intensity of the object of interest, e.g. μ and σ of $\mathrm{p}_{\mathrm{G|H}}$. In our experiments we assumed that the intensity of an object can be modelled using a normal distribution, an example of an application where this assumption can be used is shown Figure 4.4.A, where leaves are monitored using an infrared camera. However the method is not limited to normal distributions, any distribution function can be used. The normal distribution was chosen as an example since it provides good results for many applications, while only requiring two parameters to describe it. The expected intensity and standard deviation of the intensity of an object can be estimated based on the segmentation results from previous frames. Another possibility is to estimate the expected intensity based on the initialisation of the active contour and the image itself. If the initialisation is not too far from

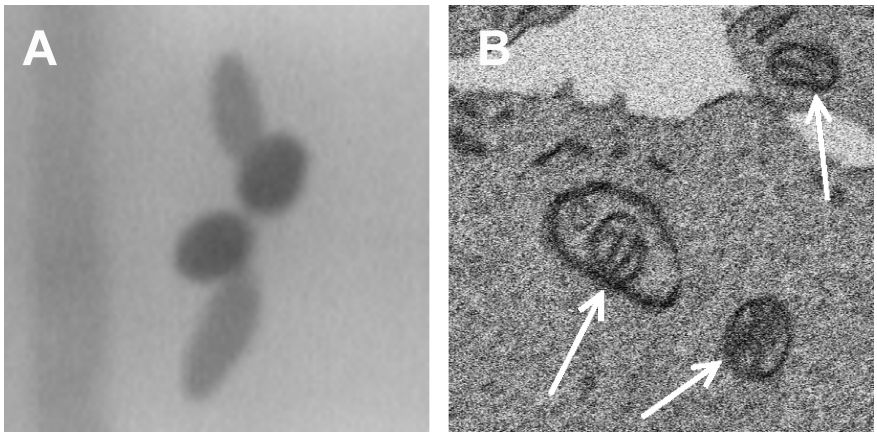


Figure 4.4: Two examples where intensity modelling can be useful for active contour segmentation. A) a thermal infrared image of plants, B) an electron microscopic image of mitochondria.

the true object, we can estimate the intensity based on the intensity enclosed by the initialisation of the active contour. To discard the influence of background or other objects, we can use the median or a rank order filter as a robust estimator of the expected intensity. Note that so far, we attempted to model the intensity for the complete object of interest, it is however only required to model the intensity along the contour. Therefore, objects with a specific intensity along the contour, but a different intensity or structure at the centre of the object, can also be segmented using the proposed method. An example of this is the analysis of a mitochondrion in an electron microscopic image. Mitochondria have a very specific texture inside, but have a dark boundary at the contour. Figure 4.4.B shows an example of such a microscopic image. White arrows indicate the mitochondria.

4.2.2 Experimental results

As a use-case, we discuss the analysis of sugar beet seedling plants in thermal infrared images. Figure 4.4.A shows an example of such a thermal image. Every hour a new image was taken. In this sequence the 4 leaves of the plant all move in different directions, at different speeds. Since this movement does not seem to have a clear motion model, we cannot incorporate any prior knowledge about the motion in our tracking methods. Hence, the segmentation result of the previous frame is chosen as the initialisation of the active contour, this without any motion compensation. Note that, for many of these seedling plants, a difference in leaf temperature can be observed. This temperature difference is mainly present

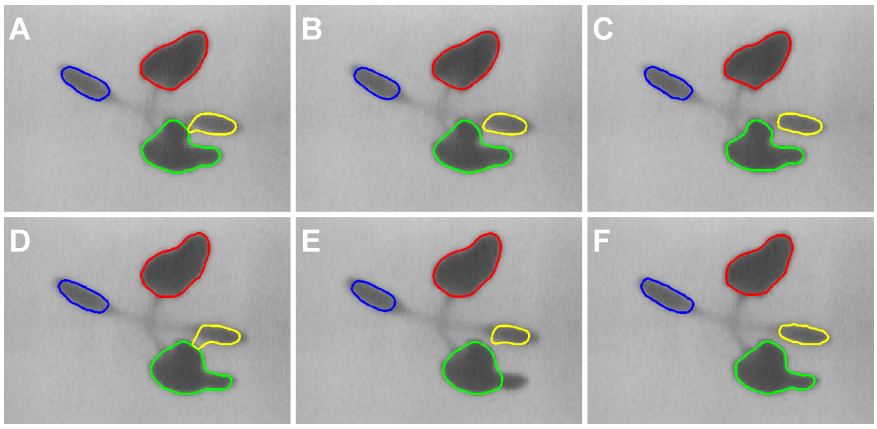


Figure 4.5: An example of segmentation and tracking using active contours. (A,B) is the segmentation result of the classical active contours, but with different parameters. (C) the segmentation result of the proposed intensity based optimisation. (D-F) shows the tracking result after processing 7 frames. (D,E), the tracking result of the classical active contours. (F) the tracking result using the proposed method [De Vylder 09a].

between older and younger leaves, i.e. between the two big leaves and the two small leaves.

The first frame is initialised using a watershed algorithm [De Bock 07]. The top row shows the segmentation result for the first frame. Fig. 4.5.A and Fig. 4.5.B show the segmentation results of the original GVF active contours, but with different weighting parameters for the internal energy. In the first example, Fig. 4.5.A, the strong attraction of the big leaf is visible. Even though the error is still small, the yellow contour is already attracted away from its true leaf boundary towards the big leaf. Fig. 4.5.B does not suffer from this problem because it puts high smoothness constraints on the shapes of the snake. It achieves this by defining high values for α and β in eq. (4.1.2). Fig. 4.5.C shows the segmentation result using IGVF. The proposed method shows a similar result as the one with high weights for the internal energy.

The improvement of the proposed method might not seem spectacular, especially compared to Fig. 4.5.B, but keep in mind that every error has the risk of propagating to subsequent frames. The second row of Fig. 4.5 shows the tracking results after analysing 7 frames. The classical active contours result in erroneous segments, as can be seen in 4.5.(D,E). In Fig. 4.5.D the error of Fig. 4.5.A has propagated. Whereas in Fig. 4.5.E the consequence of too strong shape constraint becomes prominent. As can be seen in Fig. 4.5.F, the IGVF active contour is not hampered by the presence of multiple objects, nor does it constrain the shapes of the segments.

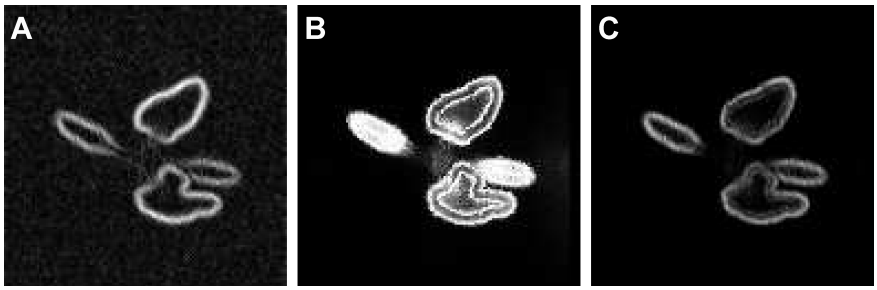


Figure 4.6: Respectively the magnitude of the gradient, the likelihood of neighbouring pixels and the proposed feature map [De Vylder 09a].

Fig. 4.6 illustrates the cause of the erroneous segmentation. Fig. 4.6.A shows the magnitude of the gradient, which is used as a feature map in the GVF force field. It is clear that the bigger leaves show a stronger gradient magnitude. It is exactly this difference in edge strength that causes the conventional active contours to fail. Fig. 4.6.B shows the likelihood that an edge pixel borders to a leaf of interest, eq. (4.2.3). The expected intensity for this likelihood was estimated for the yellow delineated leaf in Fig. 4.5.A. Note that the small elongated leaves show a high likelihood whereas the large leaves only show a high likelihood around the border. This double halo is present due to blurred edges. This blurred edges show a lower intensity than the actual leaves, and therefore the boundary has an intensity that resembles the intensity of the elongated leaves. Fortunately, these halos do not coincide with the actual edges. Fig. 4.6.C shows the final feature map, eq. (4.2.4). The halo's mentioned in Figure 4.6.B are only slightly present in this final feature map. Furthermore, the strong edges no longer dominate the feature map, hence they do not attract the active contours delineating the elongated leaves.

For the quantitative validation of the proposed method we compare the segmentation of 419 leaves with ground truth segmentation. This ground truth is the result of manual delineation of individual leaves by a botanist. As an initialisation, we started from the ground truth delineations in the previous frame. The conventional active contour resulted in an average Dice coefficient of 0.87, for an explanation of the Dice coefficient we refer to App. B. Active contours that where optimized using the proposed IGVF external force resulted in an average Dice coefficient of 0.89. The used parameters where set empirically. The exact parameters used for the active contours are $\alpha = \frac{1}{5}$, $\beta = \frac{1}{5}$, $\mu = 10$, $\eta = 2.5$ and 10 iterations used for the optimisation of the GVF and IGVF force field.

In a second experiment, we test the accuracy in a tracking application, i.e. the segmentation result of the current frame is used as initialisation for the next frame. As a test-set 16 leaves where tracked over 9 frames. Figure 4.7 shows the average

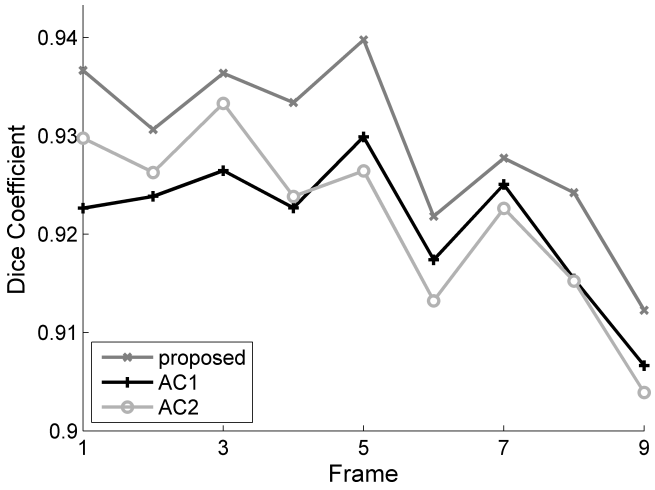


Figure 4.7: The average Dice coefficient for tracking leaves over time. The green curve shows the average Dice coefficient for an active contour optimized using the proposed IGVF. Red and Blue show the result of optimising the active contour using the original GVF. AC2 represents an active contour with high weights for the internal energy, thus strongly regularising the shape, whereas AC1 has the same weighting parameters as used for the proposed method.

Dice coefficient over time. The active contour optimized using the proposed IGVF (shown in green) is compared to the segmentation result of an active contour with the same parameters (shown in red) and an active contour with high weights for the internal energy (shown in blue), i.e. $\alpha = 1$ and $\beta = 2$. Just as in the example of Fig. 4.5, stronger regularisation does not result in more accurate tracking. The proposed method results in more accurate tracking for all frames. Note that all three methods show a strong decrease in Dice coefficient at frame 6 and 9. This is due to strong motion from the leaves at those frames. This motion results in initialisations that are further away of the actual object. For some of the leaves, the initialisation was partially outside of the capturing range of the external force, thus relying solely on the internal energy to result in accurate segmentation. This unfortunately results in less accurate segmentation, which reflects in the decrease in Dice coefficient.

4.3 Optimisation in time lapse sequences

By using an external force field, the active contour is more likely to converge to the actual object of interest, even if the active contour is not initialised in the close vicinity of the object of interest. This proper convergence is hampered by

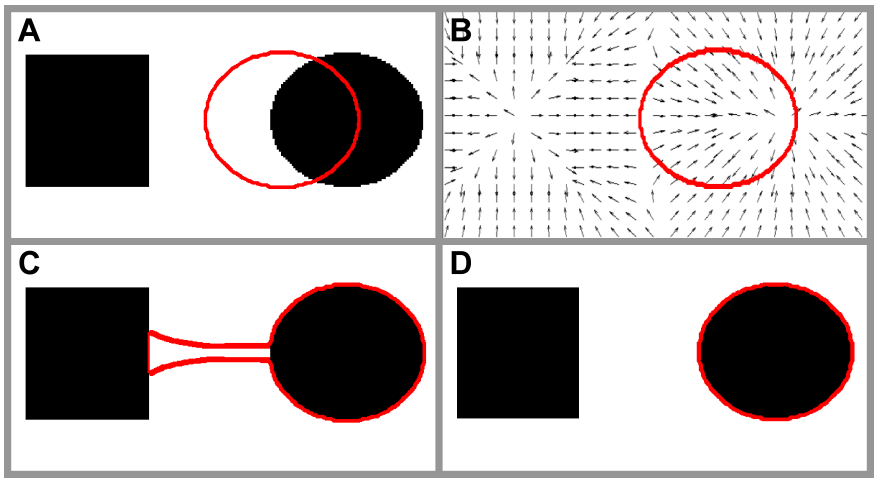


Figure 4.8: Example of segmenting the circle using a GVF optimized active contour and using the proposed optimisation. A) the initialisation of the active contour, B) the GVF force acting on the active contour, C) the segmentation result based on the GVF force field, D) the segmentation result using the proposed two-step optimisation scheme [De Vylder 09b].

the presence of clutter or similar objects. The previous subsection proposed to solve this problem by incorporating prior knowledge on the intensity of the object of interest. While this approach shows good results for objects and clutter with different intensities, it does not help in situations where similar objects are present. Figure 4.8 illustrates a simple example of this problem. Fig. 4.8.A shows the input image: two black objects on a white background. The initialisation of the active contour is shown in red. Fig.4.8.B visualizes the GVF force field. Due to the force field, even homogeneous regions display a force, guiding the active contour in a specific direction. This force field attracts the active contour to both objects, resulting in erroneous segmentation, such as can be seen in Fig. 4.8.C.

With a proper initialisation, e.g. manually annotating in semi automatic methods, this problem is less likely to occur. However, in tracking applications, the resulting active contour of frame t is used as an initialisation for the contour of frame $t + 1$ [Tsechpenakis 04]. Therefore, if the object of interest undergoes strong displacements or deformations, the risk of (partially) evolving towards the wrong object forms a serious hazard. Furthermore, an erroneous segmentation result for this time frame will be used as the initialisation of the segmentation of the next frame, thus propagating the error over time. While for many video sequences the deformations and displacements are sufficiently small between time frames, this is not the case for time lapse sequences. Time lapse sequences are low temporal frequency videos, which are often used to study slow moving biological objects,

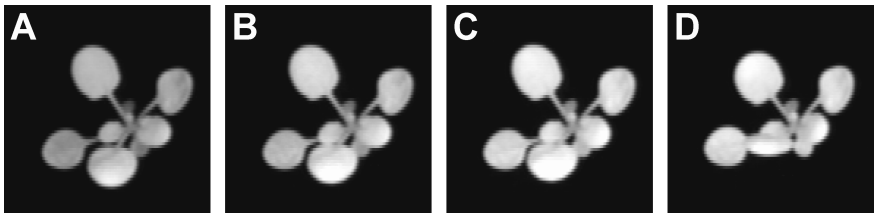


Figure 4.9: An example of four consecutive time frames of a time lapse sequence. This sequence monitors an Arabidopsis thaliana plant using a chlorophyll fluorescence camera.

e.g. to monitor plant growth.

A common approach to tackle this problem is to anticipate the movement of the object of interest [Isard 98b, Goobic 01, Ray 04, Tsechpenakis 04, Dambreville 06]. These algorithms all use some form of prior knowledge about the motion. This kind of prior knowledge is not always available and limits the number of applications. An example where object motion is hard to predict are plant leaves in time lapse sequences. Leaves undergo apparent random displacements around the stalk. An example of such drastic displacements is shown in Fig. 4.9. The first three frames show little to no displacement, whereas in the last frame the leaves in the bottom of the image suddenly move about three quarters of their own size. This leaf motion is intrinsic to the plants, so it is not possible to avoid this motion by monitoring the plants in a wind free environment. We would like to stress that these plant leaves move in an apparent random way, making it difficult to anticipate this motion. Although a motion by particle filtering approach could be used for such random displacements, it would require an excessive amount of particles, which results in a computational burden.

4.3.1 Two phase optimisation

Segmentation of time lapse sequences using active contours is prone to errors due to the low temporal resolution. This low temporal resolution causes bad active contour initialisations, which often results in erroneous segmentation. To overcome the problem of low temporal resolution, we propose to optimize the active contour in two steps: the first step moves the active contour without deforming it, while only in the second step the actual deformation is done. By optimising the active contour in two phases, the active contour is less hampered by clutter or other objects in the image. Fig. 4.8.D shows the segmentation result on the synthetic example using this optimisation approach. This first phase results in a rough detection of the object, this is done by only compensating for the displacement of the object. This step can be formulated in the active contour framework by defin-

ing a new internal energy. Before defining the actual energy, let us first define the following operator:

$$\blacktriangle \mathbf{x}(.) = \mathbf{x}(.) - \int_0^1 \mathbf{x}(s)\, ds. \tag{4.3.1}$$

The ▲ operator translates a curve such that the centre of mass of the curve coincides with the origin of the coordinate system. This operator allows us to define an internal energy function that restricts the deformation of the active contour, without limiting the displacement of the active contour. This energy function is defined as follows:

$$E_{int}[\mathbf{r}(.)] = \int_0^1 \gamma \, \|\blacktriangle \mathbf{r}(s) - \blacktriangle \mathbf{r_0}(s)\| \, ds, \tag{4.3.2}$$

where $\mathbf{r_0}(.)$ represents the initialisation of the active contour and γ is a weighting parameter. If the weight of the internal energy, γ in eq. (4.3.2) is high, i.e. γ is much larger than the average external force working on the active contour, then this optimisation results in translating the active contour as a rigid body. This new energy function can be optimized using gradient descent by treating $\mathbf{r}(s)$ as a function of time, i.e. using the update steps in eq. (4.1.5). This requires finding an $\mathbf{r}(.)$, such that

$$\frac{\partial \mathbf{r}(s,t)}{\partial t} = \int_0^1 \mathbf{v}(\mathbf{r}(s,t))\, ds \tag{4.3.3}$$

vanishes for all s. Note that $\mathbf{v}$ corresponds to one of the external force fields described in eq. (4.1.11), eq. (4.1.12) or eq. (4.2.4). Once this rigid active contour stabilizes, the pixels enclosed by the active contour can be considered to be a rough detection of the object of interest. This rough detection can be used as an initialisation for the conventional active contour deformation. So in the first step the active contour is rigidly displaced to roughly detect the object of interest, whereas in a second step, the original or IGVF active contour optimisation can be used to accurately segment the object of interest.

4.3.2 Experimental results

As an example we tracked the four leaves of a sugar beet seedling in a thermal infrared time lapse sequence. Figure 4.10.A shows an example of such a frame. The initialisation of the first frame is based on the result of a watershed algorithm [De Bock 07], see Fig. 4.10.B. Figure 4.10.C shows the tracking result using the conventional active contour optimisation after 55 frames. The conventional active contour, optimized using a GVF force is only able to track two out of the four

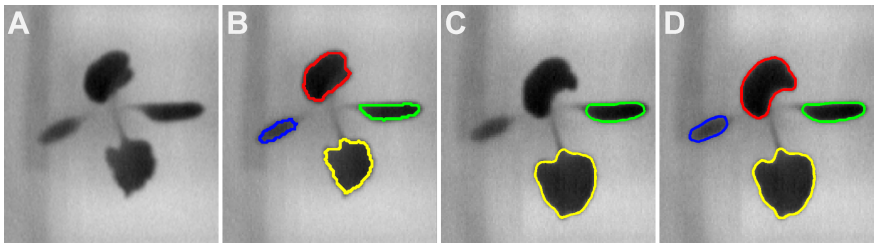


Figure 4.10: An example of tracking leaves in thermal infrared time lapse sequences. A) the first frame, B) the initialisation of the active contour in the first frame, C) the tracking result after 55 frames using conventional active contour optimisation, D) the tracking result after 55 frames using our proposed two phase optimisation [De Vylder 09b].

leaves. Figure 4.10.D shows the tracking result using our proposed two phase optimisation. This results in accurate tracking of all four leaves.

The top row of figure 4.11 shows in more detail why the conventional active contours result in erroneous tracking. Note that the most left leaf already was lost during tracking before frame 14. Between frame 14 and frame 15 one of the leaves undergoes a sudden displacement, respectively shown in Fig. 4.11.A and Fig. 4.11B. This motion results in a leaf boundary that lies almost exactly between the upper part and the bottom part of the active contour. As a consequence, the complete right part of the active contour is attracted to the same part of the leaf boundary, thus resulting in erroneous segmentation. This is not an isolated error, since this segmentation result is used as the initialisation of the active contour in the next frame. Figures 4.11.C-D show the error propagation. The active contour shrinks into a line upon reaching frame 17 (Figure 4.11.D).

The bottom row of the figure shows the tracking result by optimizing the active contour using our proposed two phase optimisation. By first anticipating on the movement, i.e. by rigidly displacing the active contour, and only deforming the active contour in a second phase, accurate segmentation is achieved. The active contour shows no problems in tracking the true boundary of the leaf over time.

For the quantitative validation, we compared the segmentation result of sugar beet seedling leaves in thermal infrared images with manual annotations. We do this for three different datasets, each containing leaves that undergo different amounts of motion. The magnitude of displacement is calculated as the Euclidean distance between the centres of mass of the leaves in the current and previous frames. The active contours are initialised by the manual annotation from the previous frame. The active contours are optimized using the following weighting parameters: $\alpha = \frac{1}{5}$, $\beta = \frac{1}{5}$, $\mu = 10$, $\eta = 2.5$. The original active contours are optimized using a GVF force, which was calculated using 10 iterations. The proposed method also uses a GVF force field of 10 iterations for the actual deformation, and for the first

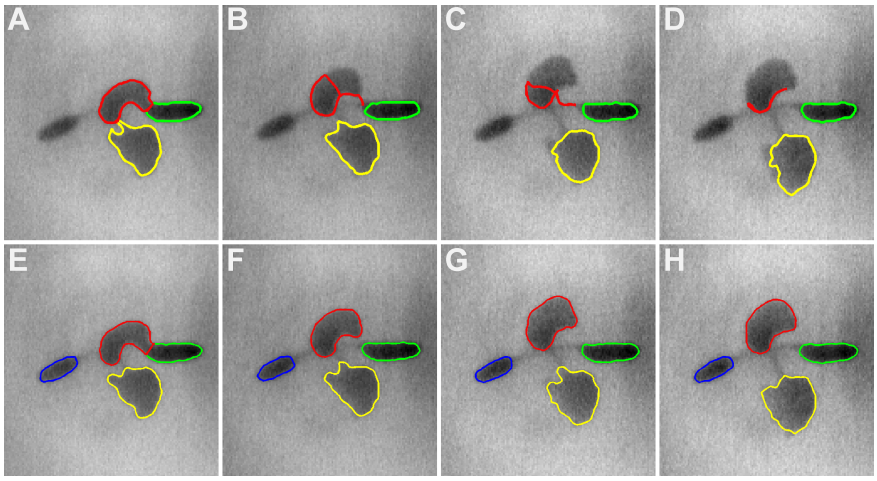


Figure 4.11: An example of tracking using active contours, frame 14 to frame 17. The top row shows the error propagation that occurs using conventional GVF optimisation. The bottom row shows the tracking result using the proposed two phase optimisation [De Vylder 09b].

phase uses a GVF force field calculated using 20 iterations. By using 20 iterations for the calculation of the GVF, we extend the capturing range of the leaves. Using 20 iterations for the GVF force used by the conventional optimisation scheme results in less accurate segmentation, because other objects further away have a tendency of improperly attracting the active contour, thus resulting in segmentation results such as in Fig. 4.8.

Table 4.1 summarises the results. "#seg." represents the number of segments in the specific dataset. Motion corresponds to the Euclidean distance (in pixels) between the centres of gravity of a leaf in the current and previous frames. Dataset1 only comprises leaves with maximum motion of 5 pixels, dataset 2 comprises leaves with a displacement between 5 and 10 pixels and the third dataset contains leaves with displacements between 10 and 15 pixels. The Dice columns show the average Dice coefficients (see app. B) for all correctly detected leaves. A leaf is considered to be correctly detected if it has a Dice coefficient higher than 0.5, otherwise it is considered as a missed segment, shown in the last two columns. For objects with little motion, i.e. dataset 1, the proposed method shows little added value, although less segments are missed. However, for larger displacements, the proposed optimisation results in significantly more accurate segmentation: higher Dice coefficients and less missed segments. Note that the concept of large displacement is relative and depends on the application. For our application, 10 pixels might seem small, but already corresponds to the width of the green and blue delineated leaves in Fig. 4.11.

		motion		Dice		missed	
	#seg.	min.	max.	original	proposed	original	proposed
dataset 1	711	0	5	**0.91**	**0.91**	41	**37**
dataset 2	94	5	10	0.83	**0.89**	33	**23**
dataset 3	31	10	15	0.71	**0.92**	20	**17**

Table 4.1: The results of segmentation using the conventional optimisation using GVF and using the proposed two phase optimisation (#seg. = number of segments in the ground truth dataset).

4.4 Optimisation using shape constraints

The discussed active contour methods help to cope with intensity differences between objects or for low temporal resolution in tracking applications. However, low contrast, clutter, blur and other image degradations can still hamper accurate active contour segmentation and tracking. An example of this problem is shown in Fig. 4.12. Figure 4.12.A shows two objects: a black square and a blurred light grey ellipse. The red ellipse represents the initialisation of the active contour for the grey object. The boundary of the black object will have a much stronger edge response than the blurred grey object, which affects the GVF force. This influence is shown in Fig. 4.12.C where due to the difference in gradient magnitude the black object shows a much wider capturing range than the grey object. This results in erroneous segmentation results, since a part of the active contour is attracted to the wrong object, see Fig. 4.12.E.

4.4.1 Shape constrained force

The gradient vector flow force tries to extend the capturing range of the active contour by assuming that the force field should be smooth except in regions where there are clear edges. This smoothness is enforced over the complete image, i.e. the capturing range is extended over the full image. This allows the active contour to be deformed towards each location in the image. However, for many applications, certain regions of the image can already be excluded. These excluded regions can be the result of an initial, rudimentary segmentation or due to some prior knowledge. An example of this is the tracking of individual leaves of plants, where the fact that all leaves are part of the plant can be exploited. While individual leaf tracking is a challenging task, the segmentation of full plants is more easily achieved (see section 3.5.2). So instead of an external force that can attract an active contour to any location in the image, we propose a new force field that attracts the active contour only within a plausible region. We not only exclude certain regions of the image as potential candidates for segments, we assume the

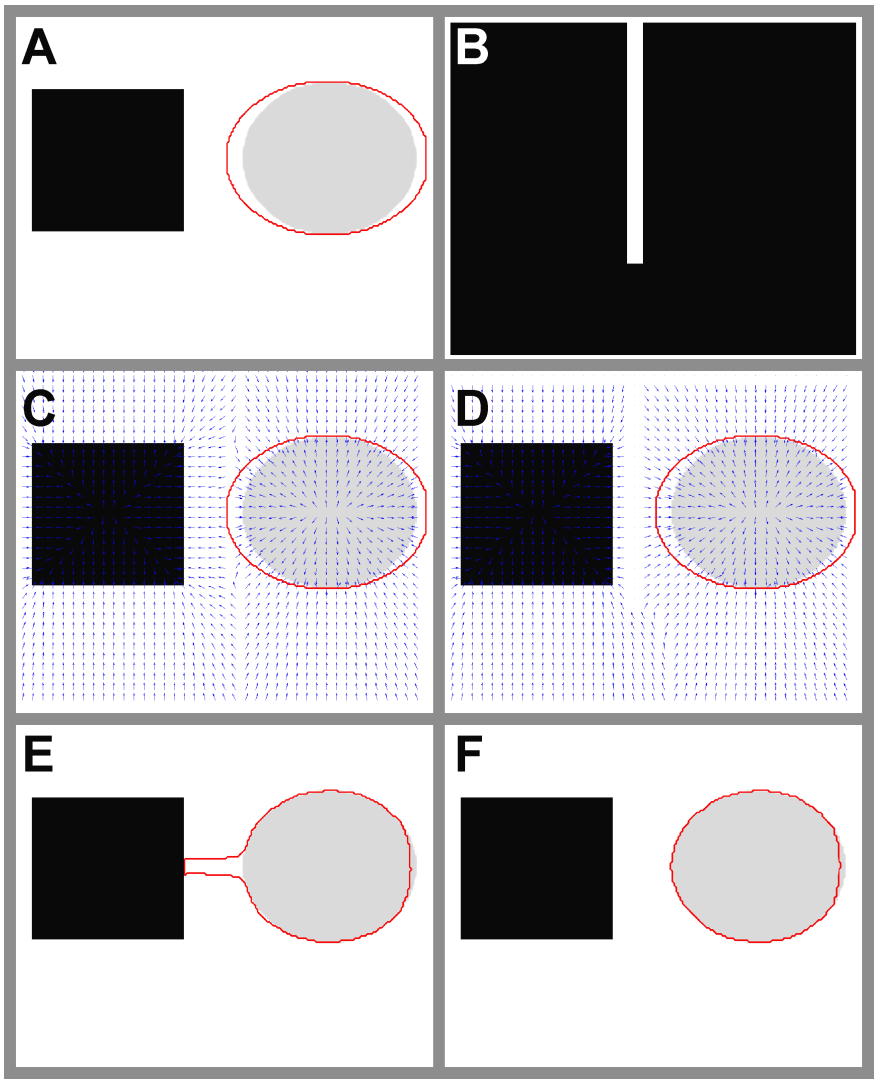


Figure 4.12: An example of active contour segmentation using the GVF force field and using the shape constrained GVF force field. A) an image containing two objects and the initialisation of an active contour (in red), B) the shape used to constrain the force field, C) the GVF force, D) the shape constrained GVF, E) the segmentation result using optimisation based on the original GVF force, F) the segmentation result using the proposed shape constrained GVF force [De Vylder 13b].

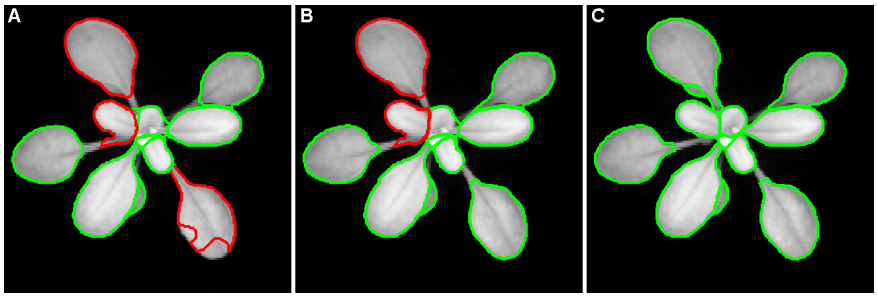


Figure 4.13: An example of segmentation using A) GVF force, B) two phase optimisation and C) the proposed SCGVF active contours [De Vylder 13b].

force field to be smooth only within a predefined shape. This new force field, $\mathbf{v}(.,.) = (u(.,.), w(.,.))$, is calculated by minimizing the following energy function:

$$E_{SCGVF}[\mathbf{v}(.,.)] = \iint_{\Omega} \mu\Big(\frac{\partial u(x,y)}{\partial x}^2 + \frac{\partial u(x,y)}{\partial y}^2 + \frac{\partial w(x,y)}{\partial x}^2 + \frac{\partial w(x,y)}{\partial y}^2\Big) + |\nabla f(x,y)|^2 |\mathbf{v}(x,y) - \nabla f(x,y)|^2 dx\, dy, \quad (4.4.1)$$

where Ω are the points belonging to the shape, i.e. all data-points that are still plausible. At points that do not belong to this shape, i.e. the background, the force is set equal to the zero vector. Just as for the GVF force, the energy function in eq. (4.4.1), consists of two terms: a smoothness term and a data fidelity term. However, now these terms only contribute to the energy for pixels belonging to the shape, Ω. Since this force field is calculated based on the GVF energy function over a shape, we will refer to this energy function as *shape constrained gradient vector flow (SCGVF)*. Figure 4.12.D shows the proposed SCGVF force field. This force field is constrained to the black shape in Figure 4.12.B. Note that this is a very rough approximation of the segments, nevertheless it will be sufficient to achieve correct segmentation. The capturing range of the black object in Fig. 4.12.D is clearly less than in Fig. 4.12.B. This is because the force only has to be smooth within the constraining shape, and as such, the force is only active within the boundary of the constraining shape. Figure 4.12.F shows the accurate segmentation result of the active contour optimized using the SCGVF force.

4.4.2 Experimental results

Fig. 4.13 shows an example of active contour tracking of individual leaves in chlorophyll fluorescence time lapse sequences. The images show the 18th frame.

The first processed frame was initialized with a manual delineation of the leaves. Fig. 4.13.A shows the tracking result of an active contour using a GVF force. This approach results in bad segmentation (shown in red), which is mainly due to the low temporal resolution. The leaves move too far for the active contour to optimize correctly. Section 4.3 proposed a two phase optimisation to overcome the challenges related to low temporal resolution. Fig. 4.13.B shows the result of this optimisation approach. The method clearly has less problems with the low temporal resolution. Nevertheless, there are still some erroneous segmentations remaining (shown in red), e.g. a small partially occluded segment is no longer found or parts of the stalk are merged with a leaf. The active contour optimised using the two phase optimisation together with the proposed SCGVF force is shown in Fig. 4.13.C. This method is able to track all leaves, also the partially occluded ones.

In order to quantitatively validate the proposed method, 12 leaves where delineated over 20 frames. The annotation of each frame was the result of a watershed segmentation, manually corrected by a human operator. The watershed algorithm missed on average 3 leaves, found 1 false segment and had 6 over or under segmentation errors for each frame. This makes it clear that the watershed algorithm might be useful in a semi-automated method, i.e. in combination with the corrections of a human operator, but is no option for real automatic analysis of leaf fluorescence time lapse sequences.

In Fig. 4.14 the average Dice coefficient (see app. B for the definition) is plotted for each time frame. The first frame was initialized by the boundaries of the ground truth segments. The proposed method results in good tracking results, whereas the regular GVF force or two phase optimisation clearly are hampered by error propagation. The error propagation of the other methods can be observed by the fast decreasing Dice coefficient starting around frame 16.

4.5 Optimisation using force propagation

By using an external force field, instead of optimizing an external energy, the active contour is less prone to erroneous optimisation. Especially the GVF force field (see eq. (4.1.11) for more details) has proven its merit. The GVF force can be used either directly or in combination with one of the previous proposed methods. The downside of GVF is that the force field requires solving a partial differential equation. This minimisation requires the iterative updating of a force field, which is both memory and time consuming. The VFC force field, another external force field (see eq. (4.1.12) for more details), is not hampered by such a computational burden. Given a specific application, it is however not straightforward to

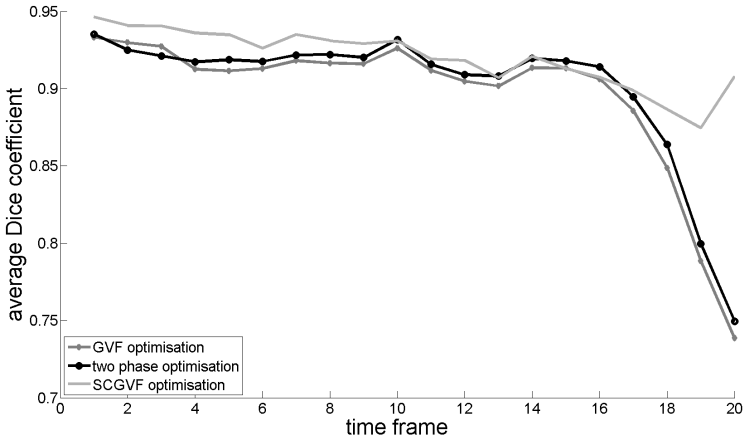


Figure 4.14: The average Dice coefficient for each frame.

choose the optimal parameters for the VFC force field. In the following section, we propose a new external energy based on edge and force propagation (*ep* and *fp* respectively). This novel force does not need the iterative optimisation of a force field, and has intuitive parameters. The proposed force field is computationally efficient to calculate, and results in faster optimisation of the active contour.

4.5.1 Two pass propagation

We start from an initial force field $\mathbf{v}_{fp}(x, y)$ which is zero everywhere and a feature map $f_{ep}(x, y)$, such as in eq. (4.1.9), where the feature map expresses how much *evidence* there is that there is a specific feature, such as an edge at pixel (x, y). The goal is to create a new feature map, where there is high edge evidence at the edge itself and where the edge evidence gradually decreases further away from the edge. The main idea is to propagate strong edge evidence at a certain pixel to its neighbouring pixels with lower edge evidence. This step only propagates edge evidence to the direct neighbouring pixels, which would also require the iterative updating of the complete force field. This can however be avoided by the following efficient dual scan line algorithm, which updates pixels solely based on their own value, and on already processed neighbours:

1. Scan the edge map row by row from top to bottom.

2. In each row, the pixels are scanned from left to right.

3. Define the already processed neighbourhood of a pixel (x,y), i.e.

$$\mathbf{n}(x,y,i) = \begin{cases} (x, y-1) & \text{if } i = 1 \\ (x-1, y-1) & \text{if } i = 2 \\ (x-1, y) & \text{if } i = 3 \\ (x-1, y+1) & \text{if } i = 4. \end{cases} \tag{4.5.1}$$

   Fig. 4.15 shows the already processed neighbour pixels for a pixel P in pink.

4. Update the pixel in the feature map by:

$$f_{ep}(x,y) = \max \begin{cases} \gamma \max_i f_{ep}\big(\mathbf{n}(x,y,i)\big) & \text{(a)} \\ f_{ep}(x,y) & \text{(b),} \end{cases} \tag{4.5.2}$$

   where $\gamma \in [0,1]$ is a weighting coefficient, which determines the speed at which edge evidence decreases. This propagates edge evidence of a pixel beyond its direct neighbours in the scanning direction. Note that all pixels of $f_{ep}(\mathbf{n}(x,y,.))$ are already processed in previous steps, due to the scanning order.

5. The feature map resulting from this scanning algorithm defines a new external energy. Instead of using the gradient of this external energy we propose to propagate the force corresponding to the external energy, i.e.

$$\mathbf{v}_{fp}(x,y) = \begin{cases} \mathbf{v}_{fp}(\mathbf{n}(x,y,i)) + \mathbf{n}(x,y,i) - (x,y) & \text{if (a) was used in eq. (4.5.2)} \\ \mathbf{v}_{fp}(x,y) & \text{if (b) was used in eq. (4.5.2)} \end{cases} \tag{4.5.3}$$

   where $i = \arg\max f_{ep}\big(\mathbf{n}(x,y,i)\big)$. Fig. 4.15.B shows an example of this force propagation. Assume that the feature map is maximal in its top neighbour, i.e. $i = 3$. The new force vector is equal to the sum the force vector in $\mathbf{n}(x,y,3)$, shown as a dashed blue arrow, and the vector pointing from the current pixel to $\mathbf{n}(x,y,3)$, shown as the small blue arrow. This results in a vector pointing from the current pixel to the same point as to which the vector in $\mathbf{n}(x,y,3)$ points, illustrated using a black arrow.

This algorithm propagates edge evidence from top to bottom and from left to right of the image. The algorithm is then repeated in order to propagate edge evidence in the remaining directions, but in opposite scanning direction, i.e. from bottom to top and from right to left. Then the already processed neighbourhood of a pixel

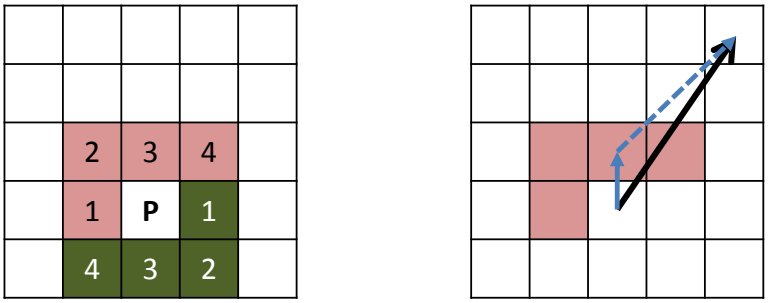


Figure 4.15: Example of edge and force propagation. A) shows the feature map, together with the two processed neighbourhoods. The pink neighbourhood is used in the first run, whereas the green neighbourhood is used for the second run. B) The force at the pixel is the sum of the force of its neighbouring pixel and the vector to that pixel [De Vylder 10e].

(x, y) is:

$$\mathbf{n}'(x, y, i) = \begin{cases} (x, y+1) & \text{if } i = 1 \\ (x+1, y+1) & \text{if } i = 2 \\ (x+1, y) & \text{if } i = 3 \\ (x+1, y-1) & \text{if } i = 4 \end{cases} \tag{4.5.4}$$

Fig. 4.15.A depicts this new neighbourhood in dark green. After the second run, the force field attracts free particles to features of interest such as edges. This force field however has a high variability in magnitude. This makes it difficult to find good weighting parameters for the regularisation terms, and as a consequence hampers the optimisation. To overcome this issue, two simple approaches can be used:

- Set the norm of a force vector equal to its value in the feature map, i.e.

$$\mathbf{v}_{fp}(x, y) = f_{ep}(x, y) \frac{\mathbf{v}_{fp}(x, y)}{\|\mathbf{v}_{fp}(x, y)\|}. \tag{4.5.5}$$

  This results in a vector field where there is a strong force in the vicinity of edges and a weak force far away of edges, i.e. far away of edges the active contour is mainly regulated by the internal energy, whereas in the vicinity of edges it is mainly deformed by the external force.

- Normalize the force:

$$\mathbf{v}_{fp}(x, y) = \frac{\mathbf{v}_{fp}(x, y)}{\|\mathbf{v}_{fp}(x, y)\|}. \tag{4.5.6}$$

  This results in an active contour that evolves at nearly constant speed, such as proposed in [Cohen 93].

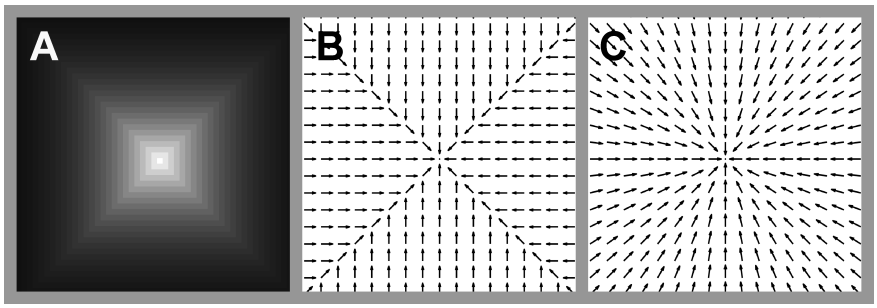


Figure 4.16: An example of edge and feature propagation for an image containing an isolated bright spot at the centre of the image. A) the feature map, resulting from the edge propagation, B) the force field defined by the gradient of the feature map, C) the force field resulting from force propagation.

Force propagation active contours can be optimized by replacing the external force in eq. (4.1.10) by $\mathbf{v}_{fp}(\mathbf{r}(s,t))$. Fig. 4.16 shows an example of the proposed force. Fig 4.16.A shows the energy map, $f_{fp}(x,y)$, of an image with an isolated point in the centre of the image. The energy map intensity gradually increases towards the centre of the image. Fig 4.16.B shows the gradient of the feature map, whereas in Fig 4.16.C the proposed force field can be seen. Note that a free particle placed in both force fields will move to the same point, i.e. the maximum of the feature map. The particle moving according to the proposed force field will move the shortest distance possible to arrive at the steady state, which is clearly not the case for the particle moving according to the gradient of the feature map.

4.5.2 Parameter settings

The desired capturing range of an active contour depends on the application. For applications where the object of interest has strong edges, the active contour should be attracted towards these strong edges, even if the active contour is in the vicinity of a weak edge. This is however an undesired property for applications where the object of interest has (partially) low contrast borders. In Fig. 4.17 we give an example of how the proposed external force field can easily be adapted depending on the application. This figure shows the proposed force field for three different parameter settings. This force field is calculated for an image containing two edges. The edge on the right of the image has a gradient magnitude that is 4 times higher as the edge on the left. The capturing range of an edge is easily tuned by changing the remembering factor, γ in eq. (4.5.2). Fig. 4.17.A shows the external force calculated with γ equal to 0.7. The magnitude of the force is proportional to the value in the energy map, $f_{fp}(\mathbf{r}(s,t))$, and is shown by the colour of the arrows. As can be seen, the magnitude of the force is much higher in the vicinity of an edge

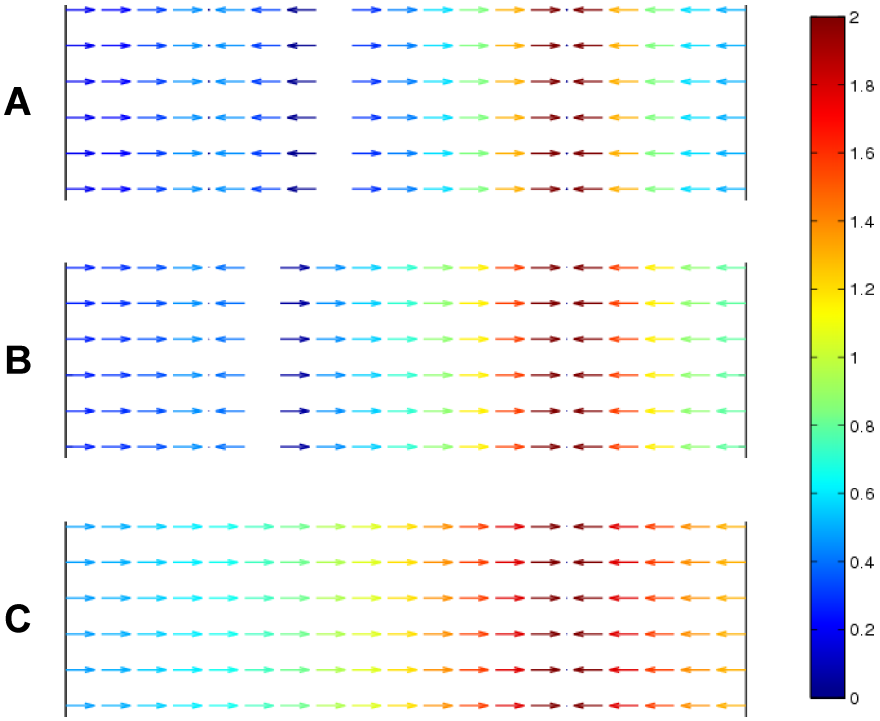


Figure 4.17: An example of the proposed external force using different parameter settings.

than for pixels further away from an edge. The weak edge captures about one third of the image between the two edges. In Fig. 4.17.B the force field is shown for γ equal 0.8. Only in the close vicinity of the weak edge would an active contour be attracted to the left edge. In Fig. 4.17.C the weak edge is completely discarded and every point of the image has a force pointing towards the right edge. This is achieved by setting γ to 0.9.

4.5.3 Directional force

While external forces significantly extend the capturing range of active contours, they do remain sensitive to some extent to the initialisation. We already proposed an optimisation strategy to overcome bad initialisations, see section 4.3, but sometimes the initialisation is too far off the actual object to result in correct segmentation. Fig. 4.18.B shows an example of such a situation. In this example, a cell nucleus moves from the centre of the image to the lower right corner. The initialisation is the result of the active contour in the previous frame, hence the initialisation (illustrated by a red dashed line) is still in the central part of the image.

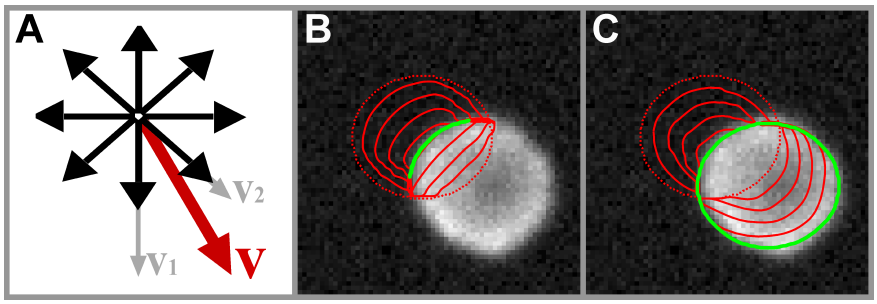


Figure 4.18: An example of directional favoured forces. A) The decomposition of the direction to favour, into the eight basic directions. B) An example of active contour segmentation using the GVF force field. C) The segmentation result using the proposed force field with prior knowledge on the direction of the motion of the object.

Although the active contour is optimised using a GVF force, the active contour is not able to segment the entire cell nucleus. Instead, the contour shrinks to only a small part of the border of the cell, shown in green. This problem is addressed in [Kalman 60, Isard 98b, Ray 04, Dambreville 06, Tang 08], where prior knowledge about the movement of the object is used. With prior knowledge on both direction and magnitude of the displacement, the initialisation can be updated accordingly [Kalman 60, Isard 98b, Dambreville 06]. For applications where only the motion direction is known, correct segmentation can be achieved by manipulating the external force field [Ray 04, Tang 08]. The main idea in these techniques is that a certain direction is favoured in the force field. The GVF force field can be adjusted to favour a certain direction by using different weights on the smoothness constraints in eq. (4.1.11). This same goal is possible with the VFC force field, by using a non-symmetrical kernel in eq. (4.1.12). The concept of favouring a specific direction can be adopted in our proposed framework of edge and force propagation by using a different remembering factor, γ_i in eq. (4.5.2), for each possible direction. The different remembering factors are calculated using the following steps:

1. Define the direction that should be favoured by the vector $\mathbf{v}$. The magnitude of this vector defines the extra weight this direction should get.

2. The proposed framework propagates edges in 8 directions, whereas $\mathbf{v}$ can be any direction. Let us define the eight possible directions by $\mathbf{d_i}$ (shown in black in Fig. 4.18.A). These direction vectors are normalised, such that they have a magnitude equal to one. The vector $\mathbf{v}$ is enclosed by two directional vectors $\mathbf{d_i}$. In Fig. 4.18.A the favouring direction vector $\mathbf{v}$, shown in red, is enclosed by the *down* and *down-right* vectors, respectively $\mathbf{d_1}$ and $\mathbf{d_2}$. We project $\mathbf{v}$ onto $\mathbf{d_1}$ and $\mathbf{d_2}$, resulting in vectors $\mathbf{v_1}$ and $\mathbf{v_2}$. The corresponding

direction vectors, $\mathbf{d_i}$, are updated by adding respectively the vectors $\mathbf{v_1}$ and $\mathbf{v_2}$.

3. Finally the direction remembering factors are defined by $\gamma_i = \gamma \frac{|\mathbf{d_i}|}{1+|\mathbf{v}|}$, where γ defines the original (directional invariant) remembering factor from eq. (4.5.2). Note that $1+|\mathbf{v}|$ is the maximal magnitude that any $\mathbf{d_i}$ can have, this corresponds to the situation that $\mathbf{v}$ is collinear with one of the direction vectors $\mathbf{d_i}$. Hence, the normalisation by $1 + |\mathbf{v}|$ ensures that γ_i is smaller than one, such that edges with high edge evidence are not overruled by weaker neighbouring edges.

Fig. 4.18.C shows that by incorporating prior knowledge about the movement into our framework, the active contour is able to segment the cell nucleus.

4.5.4 Experimental results

Convergence To test the active contour convergence using the proposed external force, a database with leaf shapes was used. This test set was chosen because it contains realistic shapes, i.e. showing both concavities and protrusions. For more details on the leaf database we refer to App. A. The active contour was initialised by a square delineating the full image. The Dice coefficient between the active contour and the leaf was measured every 10 iterations. Fig. 4.19 shows the average convergence of the active contours. The proposed force (FP) is compared with the GVF and VFC forces [Xu 98, Li 07a]. The VFC force field was calculated using a kernel based on a quadratic function as proposed in [Li 07a], the GVF force fields was calculated with $\mu = 0.2$. The proposed method converges approximately to the same result as active contours optimised using GVF, but converges significantly faster. It reaches a Dice coefficient of 0.88 in 40 iterations, compared to optimisation using a GVF force, which needs approximately 48 iterations to achieve the same Dice coefficient. Active contours optimized using VFC evolve slower than active contours using the proposed force propagation. Furthermore, they do not reach the same accuracy as active contours optimised using GVF or the proposed force propagation method.

In a second experiment, the proposed technique is tested on a microscopic dataset. The goal is to segment isolated cells in fluorescent micrographs. An example of such an image is shown in Fig. 4.18.B. Twenty cells were manually annotated by the author and compared with the results from active contours optimised using both the proposed and the GVF force field. For both methods, experiments were done with several parameter combinations, i.e. μ in eq. (4.1.11) and γ in eq. (4.5.2). For GVF, $\mu = 0.22$ resulted in the best average dice coefficient, and $\gamma = 0.97$ gave the best result for the proposed method. The resulting Dice coefficients for these

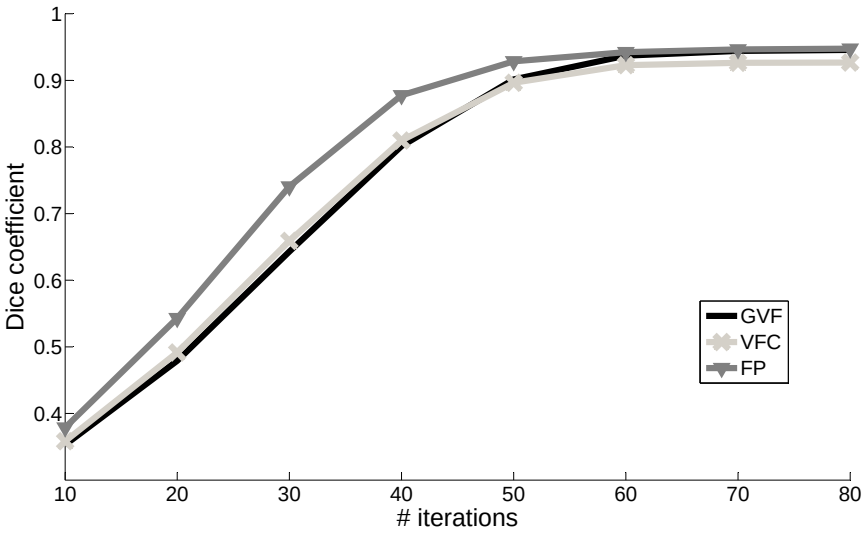


Figure 4.19: The convergence of active contours optimized using different force fields.

optimal parameters are compared in Fig. 4.20. Both methods perform adequately, with almost all Dice coefficients between 0.8 and 1, the big exception being cell 18, where the active contour optimised using GVF converged to a false optimum due to clutter in the micrograph.

Computational cost The proposed scanning algorithm recalculates the force of each pixel twice, resulting in a $O(N^2)$ algorithm for a square image with dimension $N \times N$. The VFC force computation has a complexity of $O(N^2 \log N^2)$ which is determined by the complexity of the 2D FFT and IFFT algorithms used. Since the GVF field needs $O(N^2)$ operations for each iteration and N iterations are usually needed to calculate the optimal force field [Li 07a], the GVF field has an $O(N^3)$ complexity. In Fig. 4.21 the computation time of GVF, VFC and the proposed force field are compared in function of the image size. Note that the time axis is log scaled. These experimental results[1] were calculated on a computer with an Intel i7 Q720 1.6 GHz CPU and 4 GB RAM. In agreement with the theoretical complexity analysis, the GVF field is the slowest to calculate. The VFC field is much faster than GVF, but is significantly slower than the proposed method, while the proposed method outperforms VFC on segmentation quality as well.

[1] All algorithms were programmed in C.

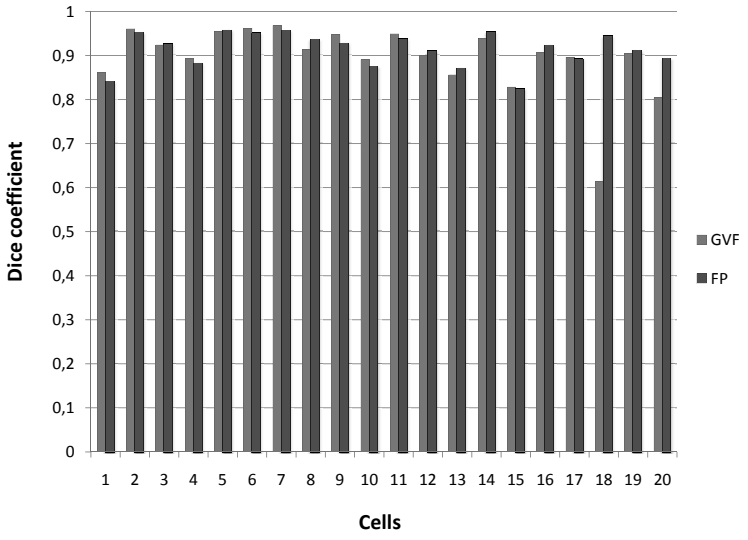


Figure 4.20: The resulting Dice coefficients of active contours optimized using GVF force and the proposed force propagation.

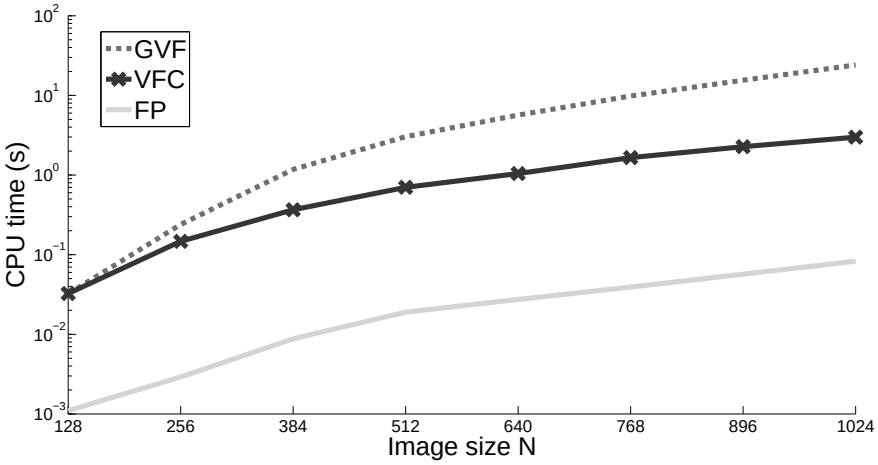


Figure 4.21: Computational cost of GVF, VFC and force propagation for a NxN image.

4.6 Probabilistic active contours

In order to guide the active contour towards the object of interest, without being hampered by clutter and noise, one or more regularisation terms are incorporated in the energy function. The influence of these regularisation terms can be adjusted by a set of weighting parameters. This enables the active contour framework to adjust to specific segmentation applications. The optimal parameters are a trade-off, where the regularisation weights have to be set high enough to overcome the influence of clutter and low enough to accurately detect the true object boundary. The tuning of weighting parameters is generally done by trial and error. This is a time consuming and error prone approach. Even after manually tuning, the parameters might not be optimal, i.e. the active contours might result in more accurate segmentation with other parameters. Note that it is non-trivial to express the segmentation accuracy in function of the weighting parameters. So finding the real optimal parameters requires exploring the full parameter space. In this section we propose a new active contour framework based on probability theory. Instead of exhaustively searching optimal weighting parameters, the proposed method uses prior knowledge about the probability of certain features, e.g. edges. It also removes the linear influence of image features, i.e. it is not because the gradient magnitude of an edge is twice as high as another edge, that it is twice as likely to be the true border of an object. This is especially relevant for segmentation applications with noisy or blurred images containing multiple objects.

The goal of our active contour framework is to find the contour, $\mathbf{r}^*(.)$, which is most likely to delineate the object of interest. This can be formalized as finding the contour that maximizes $\mathrm{p}\big(H(\mathbf{r}(.))\big)$, where $H(\mathbf{r}(.))$ is the hypothesis that the contour $\mathbf{r}(.)$ delineates the object of interest. Lets assume that in order to find such a contour we can use a set of image features $f(x, y)$, e.g. the edge strength at a specific pixel. The optimal contour is then defined as

$$\mathbf{r}^*(.) = \arg\max_{\mathbf{r}(.)} \mathrm{p}\big(H(\mathbf{r}(.))\big|f(.,.)\big). \tag{4.6.1}$$

Using Bayes' rule, we can rewrite this as

$$\begin{aligned} \mathbf{r}^*(.) &= \arg\max_{\mathbf{r}(.)} \frac{\mathrm{p}\big(f(.,.)\big|H(\mathbf{r}(.))\big)\,\mathrm{p}\big(H(\mathbf{r}(.))\big)}{\mathrm{p}\big(f(.,.)\big)} \\ &= \arg\max_{\mathbf{r}(.)} \left(\log \frac{\mathrm{p}\big(f(.,.)\big|H(\mathbf{r}(.))\big)}{\mathrm{p}\big(f(.,.)\big)} + \log \mathrm{p}\big(H(\mathbf{r}(.))\big) \right). \end{aligned} \tag{4.6.2}$$

Equivalent to the active contour energy, we can split this part in a regularisation term and a data term. To emphasise this similarity we will use similar termino-

logy, i.e. an internal and external probability, which are defines as: $\mathrm{p}_{int}(\mathbf{r}(.)) = \log \mathrm{p}\big(H(\mathbf{r}(.))\big)$ and $\mathrm{p}_{ext}(\mathbf{r}(.)) = \log \frac{\mathrm{p}\big(f(.,.)\big|H(\mathbf{r}(.))\big)}{\mathrm{p}\big(f(.,.)\big)}$.

4.6.1 Internal probability

The internal probability is completely independent of the image and can be used to incorporate the shape likelihood of an object of interest. As an example, we will use a simple model proposed in [Xu 00], which only depends on the second derivative of the contour. This results in the following internal probability:

$$\begin{aligned} \mathrm{p}_{int}(\mathbf{r}(.)) &= \log \mathrm{p}\left(H(\mathbf{r}(.))\right) \\ &= \log \mathrm{p}\left(\left|\frac{\partial^2 \mathbf{r}(s)}{\partial s^2}\right|\right) \\ &= \int_0^1 \log \mathrm{p}\left(\left|\frac{\partial \mathbf{r}^2(t)}{\partial t^2}\right|\right) dt. \end{aligned} \tag{4.6.3}$$

For this last step we assume that the second derivative of $\mathbf{r}(t)$ is independent for every t. The probability distribution $\mathrm{p}\left(\left|\frac{\partial \mathbf{r}^2(t)}{\partial t^2}\right|\right)$ can be estimated using a small training set of ground truth segments. Note that this is just an example of an internal probability model. If the objects of interest have a specific shape or if they have more pronounced local features, e.g. jags, then the internal probability could be formulated using a more complex shape model, such as the shape models proposed in [Staib 92, Rochery 06, Charmi 08].

4.6.2 External probability

The external probability depends on the used image features. These features are used to characterize an object. As an example we will model the objects of interest using an edge map, e.g. $f(x,y) = |\nabla g(x,y)|$. Under the assumption that the gradient is independent for all (x,y), the external probability can be rewritten as:

$$\begin{aligned} \mathrm{p}_{ext}(\mathbf{r}(.)) &= \log \frac{\mathrm{p}\left(f(.,.)\big|H(\mathbf{r}(.))\right)}{\mathrm{p}\left(f(.,.)\right)} \\ &= \log \frac{\mathrm{p}\left(|\nabla g(\mathbf{r}(.,.))|\big|H(\mathbf{r}(.))\right)}{\mathrm{p}\left(|\nabla g(\mathbf{r}(.,.))|\right)} \\ &= \iint_{\Omega^+} \log \frac{\mathrm{p}\big(|\nabla g(\mathbf{r}(u,w))|\ \big|H(\mathbf{r}(.))\big)}{\mathrm{p}\big(|\nabla g(\mathbf{r}(u,w))|\big)} du\, dw \\ &\quad + \iint_{\Omega^-} \log \frac{\mathrm{p}\big(|\nabla g(\mathbf{r}(u,w))|\ \big|H(\mathbf{r}(.))\big)}{\mathrm{p}\big(|\nabla g(\mathbf{r}(u,w))|\big)} du\, dw, \end{aligned} \tag{4.6.4}$$

where $\Omega^+ = \{(u,w)\big|\big(\exists k \in [0,1]\big)\big(\mathbf{r}(k) = (u,w)\big)\}$ and $\Omega^- = \{(u,w)\big|\big(\nexists k \in [0,1]\big)\big(\mathbf{r}(k) = (u,w)\big)\}$. So Ω^+ is the set of points that lie on the contour of the object of interest, whereas Ω^- contains all other points. For many biological monitoring applications, imaging happens in a strictly controlled environment. Due to this controlled setup, technicians can avoid clutter. By extension, they can avoid objects with the same type of edges as objects of interest. Therefore, we assume $\iint\limits_{\Omega^-} \log \frac{\mathrm{p}\left(|\nabla g(\mathbf{r}(u,w))| \ \middle| H(\mathbf{r}(.))\right)}{\mathrm{p}(|\nabla g(\mathbf{r}(u,w))|)} du\ dw$ to be small. This allows us to approximate eq. (4.6.4) by discarding this term, i.e.

$$P_{ext}(\mathbf{r}(.)) = \int_0^1 \log \frac{\mathrm{p}(|\nabla g(\mathbf{r}(t))| \ \big|H(\mathbf{r}(.))\big)}{\mathrm{p}(|\nabla g(\mathbf{r}(t))|)} dt. \tag{4.6.5}$$

The probability densities used in eq. (4.6.5) can be interpreted as follows:

- $\mathrm{p}(|\nabla g(\mathbf{r}(t))| \ \big|H(\mathbf{r}(.))\big)$: the probability density that the gradient of a point lying on the contour of a real segment is equal to $|\nabla g(\mathbf{r}(t))|$.
- $\mathrm{p}(|\nabla g(\mathbf{r}(t))|)$: the probability density that the gradient of a random point in the image is equal to $|\nabla g(\mathbf{r}(t))|$, whether this point is on the contour of a real object or not.

The probability distribution of the gradient magnitude of a pixel at the object boundary can be estimated from a small training set of images where the objects are manually delineated: measure the gradient magnitude at the contours delineating the ground truth segments. Based on these measurements, calculate the probability distribution, e.g. using a kernel density estimator. The probability distribution of the gradient magnitude, irrespective of being located at the contour, can be estimated in a similar way, but instead of measuring only the gradient strength at the contours, measure it at each pixel in the training data set. Note that although this example uses the gradient, this framework could also incorporate other image features such as ridges, intensity, region-based features, etc. [Blake 97, Poon 97, Chan 99].

4.6.3 Optimisation

Substituting the proposed internal and external probabilities in eq. (4.6.2) results in:

$$\begin{aligned} \mathbf{r}^*(.) &= \arg\max_{\mathbf{r}(.)} \Big(\mathrm{p}_{int}(\mathbf{r}(.)) + \mathrm{p}_{ext}(\mathbf{r}(.)) \Big) \\ &= \arg\max_{\mathbf{r}(.)} \int_0^1 \log \frac{\mathrm{p}(|\nabla g(\mathbf{r}(s))| \ \big|H(\mathbf{r}(.)])\big)}{\mathrm{p}(|\nabla g(\mathbf{r}(s))|)} + \log \mathrm{p}\left(\left| \frac{\partial \mathbf{r}^2(s)}{\partial s^2} \right| \right) ds. \end{aligned} \tag{4.6.6}$$

This can be optimised by solving the following Euler-Lagrange equations:

$$\begin{aligned} \frac{dx(s,t)}{dt} &= \frac{d \log \mathrm{p}\left(\left|\frac{\partial \mathbf{r}^2(s,t)}{\partial s^2}\right|\right)}{ds} + \frac{\partial \log \frac{\mathrm{p}\left(|\nabla g(\mathbf{r}(s,t))| \,\middle|\, H(\mathbf{r}(.,t))\right)}{\mathrm{p}(|\nabla g(\mathbf{r}(s,t))|)}}{\partial x} \\ \frac{dy(s,t)}{dt} &= \frac{d \log \mathrm{p}\left(\left|\frac{\partial \mathbf{r}^2(s,t)}{\partial s^2}\right|\right)}{ds} + \frac{\partial \log \frac{\mathrm{p}\left(|\nabla g(\mathbf{r}(s,t))| \,\middle|\, H(\mathbf{r}(.,t))\right)}{\mathrm{p}(|\nabla g(\mathbf{r}(s,t))|)}}{\partial y}. \end{aligned} \quad (4.6.7)$$

In order to solve this using gradient descent, the initial $\mathbf{r}(s)$ should be in the vicinity of the true object. To avoid the probabilistic active contour of evolving to an erroneous segment, the same strategy as for conventional active contours can be used, i.e. using an external force such as GVF. This is easily achieved by using the following function as a feature map for the calculation of an external force: $f(x,y) = \log \frac{\mathrm{p}\left(|\nabla g(x,y)| \,\middle|\, H'(x,y)\right)}{P(|\nabla g(x,y)|)}$, where $H'(x,y)$ represents the hypothesis that (x,y) lies on the contour of an object of interest.

4.6.4 Experimental results

Fig. 4.22 shows an example of cell nuclei segmentation in fluorescence microscopic images using active contours. Fig. 4.22.A shows the initialisation used by the active contours. Fig. 4.22.B shows the segmentation result using a conventional active contour. The parameters where empirically set. The active contour is attracted to parts of the image that have high contrast. In general, these high contrast parts correspond to edges delineating the object of interest. This image however also shows high intensity blobs, corresponding to subcellular particles such as telomeres. The difference in intensity between these blobs and the rest of the nucleus is higher than the difference in intensity between the nucleus and the background. As a consequence, the active contour is attracted towards these blobs instead of towards the nuclear boundary. By modelling the appearance of the nuclear edges, the proposed probabilistic active contour is not hampered by these bright objects. Fig. 4.22.C shows the segmentation result of the proposed method. The active contour follows the true nuclear boundary irrespective of the vicinity of subcellular particles.

In a second example we analyse sugar beet seedlings in thermal infrared images. These thermal images are noisy, low contrast grey scale images. Fig. 4.23 shows an example of leaf segmentation using both the conventional active contours as our proposed probabilistic active contours. All active contours are using a GVF force. The GVF force was calculated using 30 iterations and a smoothing factor μ equal to 0.1. In Fig. 4.23.A the initialisation of the four different active contours is shown. This initialisation is already a good approximation of the real leaf

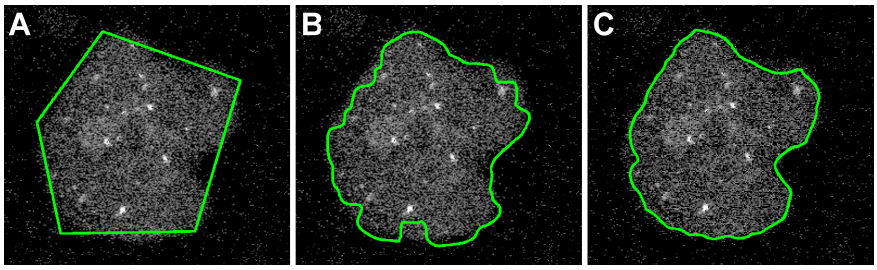


Figure 4.22: An example of segmentation of a cell nucleus in a fluorescence microscopic image. A) a rough manual delineation of the nucleus, this serves as the initialisation of the active contour. B) the segmentation result of a conventional active contour. C) the segmentation result using the proposed probabilistic active contours.

contours. This proper initialisation is however insufficient for the conventional active contours to evolve towards the real leaf boundaries. Two examples of active contour segmentation using different weighting parameters, α and β in eq. (4.1.2), can be seen in Fig. 4.23.B and Fig. 4.23.C. The active contours in Fig. 4.23.B were optimized using $\frac{1}{6}$ and $\frac{1}{3}$ respectively for weighting parameters α and β. These weighting values are apparently too low to prevent the contours to evolve to erroneous segments. An example of such an incorrect convergence can be seen at the yellow contour, which partially converged towards the border of a wrong leaf. The cause of these segmentation errors is the difference in the gradient strength. Fig. 4.24.A shows the gradient magnitude in the image. The bigger leaves display a much stronger gradient, which attract the contours of the elongated smaller leaves. Note this effect near the stem of the yellow delineated leaf. In Fig. 4.23.C., the yellow contour converged correctly by increasing the contour weighting parameters, i.e. values $\frac{5}{3}$ and $\frac{2}{3}$ respectively for parameters α and β. However due to these strong smoothness constraints, the green and blue contours lose the actual leaf borders near the tip of the leaves. Clearly, a general set of weighting parameter values is difficult to find. Even if such optimal parameter combination could be found for one image, it is not guaranteed that they would result in accurate segmentation for the full sequence.

Fig. 4.23.D shows the resulting segmentation using our proposed method. In this example the proposed probabilistic active contours do not suffer the linear influence that conventional active contours have. In order to use the proposed method, the internal and external probabilities have to be modelled first. The prior probabilities used for our method were learned using a single ground truth image. This image originated from a training time lapse sequence that was manually segmented. The distributions were calculated using the kernel density estimator with a normal distribution as kernel. In Fig. 4.24.B the external probability at each pixel is shown. The gradient strength at leaf borders varies between leaves, nevertheless

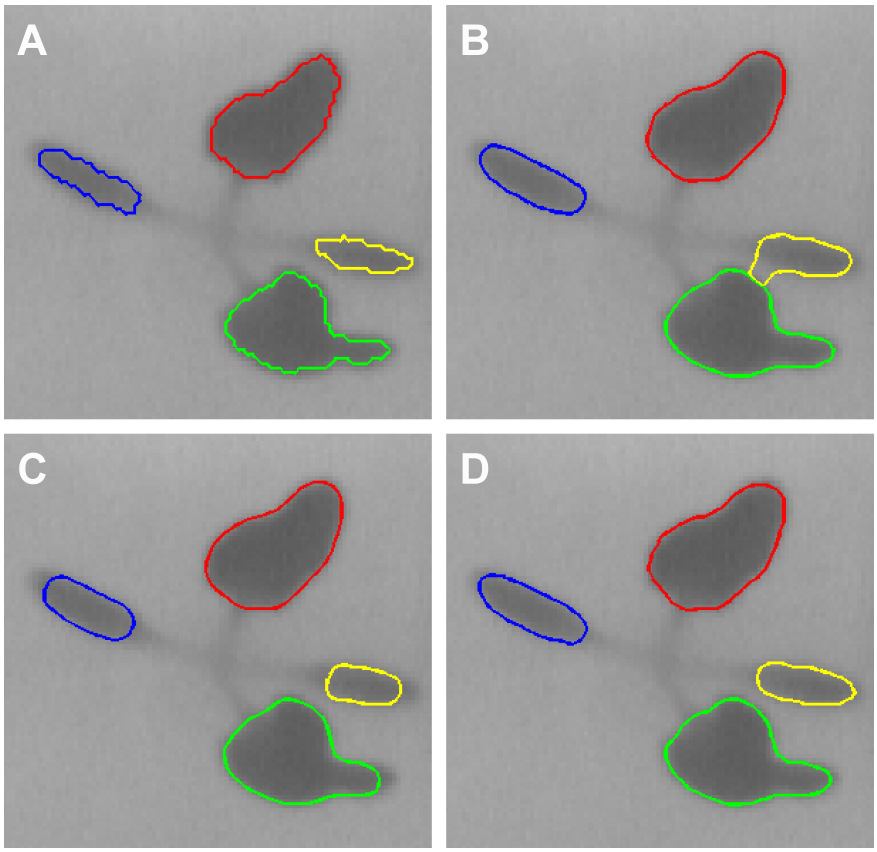


Figure 4.23: Examples of segmentation using (probabilistic) active contours. A) the initialisation used for the (probabilistic) active contours, B) the segmentation result of the classical active contours with $\alpha = \frac{1}{6}$, $\beta = \frac{1}{3}$, C) the segmentation result of the classical active contours with $\alpha = 1$, $\beta = \frac{2}{3}$, D) the segmentation result of the proposed probabilistic active contours [De Vylder 11d].

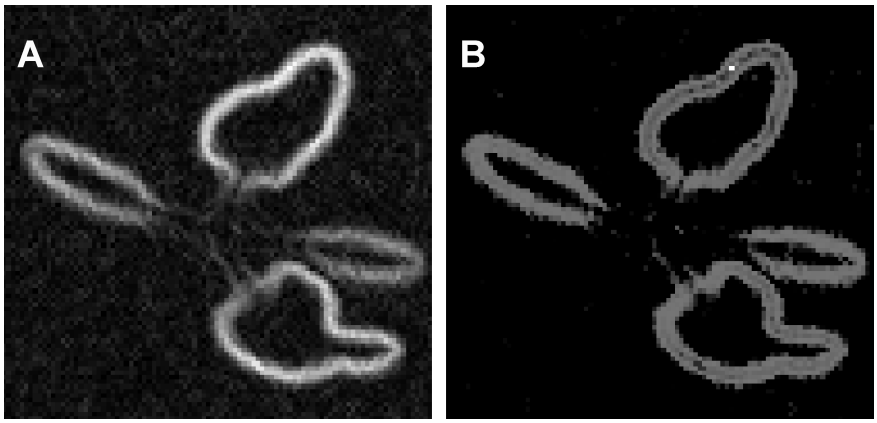


Figure 4.24: the gradient strength of the image, (e) the external probability of each pixel in the image [De Vylder 11d].

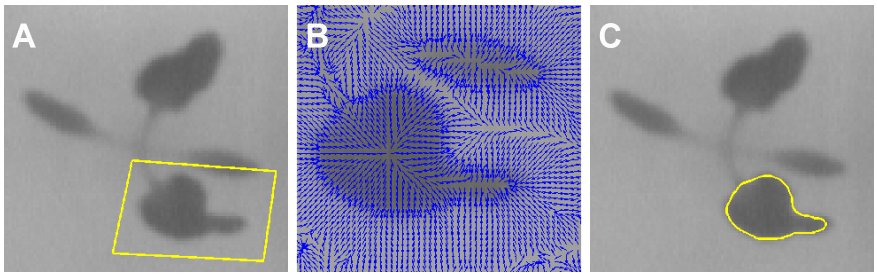


Figure 4.25: An example of segmentation using probabilistic active contours. A) the initialisation used for the (probabilistic) active contours, B) the gradient vector flow of the external probability, i.e. the force used to optimize the external probability of the active contour C) the segmentation result of the proposed probabilistic active contour [De Vylder 11d].

the different leaves show an equally strong probability at their border. This results in a better segmentation result as can be seen in Figure 4.23.D.

The previous example started from an almost perfect initialisation. This was helpful to illustrate the problems that might occur with conventional active contours, but it is rare that such good initialisations are available. Fig. 4.25 shows a more realistic example, where a rough initialisation is manually selected. Figure 4.25.A illustrates the initialisation of the active contour; whereas Fig. 4.25.B shows a detailed view of the external force field based on the external probability. The force field points, as desired, towards both leaves. The regularisation effect of the internal probability enforces the contour to converge towards the correct leaf, see Fig. 4.25.C.

As a last example, our proposed method was applied for leaf tracking in a time

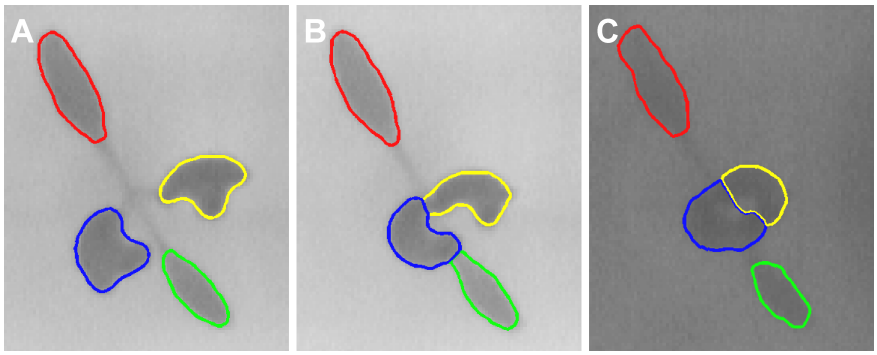


Figure 4.26: An example of tracking using probabilistic active contours. These show the tracking result of respectively time frames 1, 7 and 15 [De Vylder 11d].

lapse sequence. As can be seen in Fig. 4.26 the proposed method copes with the movement and deformation of the leaves. Even frame 15 where the illumination level diminished due to nightfall, is still segmented correctly. If this illumination change is too strong, the learned probability distributions will not correspond to the image features. Therefore, the segmentation results will be less accurate. This can be seen in Fig. 4.26.C, where the contours delineating the bigger leaves slightly miss the true border at the centre of the plant. Although this error is almost unnoticeable at this frame, there is a risk that it becomes more prominent in subsequent frames that will use these contours as an initialisation. This problem could be resolved by updating the internal and external probabilities based on the previous frames [Elhabian 08, Borghgraef 10].

To quantitatively validate the proposed method, 56 images each containing 4 leaves have been manually annotated by a botanist. This dataset has the same illumination changes as mentioned for Fig. 4.26. The initialisation of the active contours was based on these ground truth segments: the segments were dilated using a circular structuring element of size 5, the borders of these dilated segments were used as initialisation. As a validation metric the Dice coefficient is used (see Appendix B). In order to compare our method with the active contours with the most optimal parameter setting, the image sequence has been segmented using different parameter settings, i.e. α and β both in the range of $0, \frac{1}{30}, \frac{2}{30}, \frac{3}{30}, ..., 1$. The best combination of α and β resulted in an average Dice coefficient of 0.872. Using these optimal parameters, 24 segments resulted in a Dice coefficient equal to 0, i.e. the segments where completely lost. The proposed probabilistic active contours resulted in an average Dice coefficient of 0.929 with no leaves lost.

4.7 Conclusion

This chapter describes the popular parametric active contour segmentation and tracking framework. The chapter gave a brief overview of the state-of-the-art of this framework. Furthermore, we proposed several new contributions to the parametric active contour framework:

- We proposed a new external force, which in contrast to other state-of-the-art forces does not only use edge information but also incorporates prior knowledge of the intensity of the object of interest. This helps to segment individual objects in an image that contains multiple objects, each with different average intensities. This new force field was used for the segmentation of individual sugar beet seedling leaves in thermal infrared images. The proposed method resulted in an average Dice coefficient of 0.89, compared to a Dice coefficient of 0.87, which is achieved by the conventional GVF force.

- We discussed how the active contour framework could be used for the analysis of time lapse sequences. The low temporal resolution of these sequences normally causes severe problems for conventional active contour tracking. However, our proposed two phase optimisation scheme allows accurate tracking even for low temporal sequences. This is achieved by first optimising the displacement of the active contour, and only in a second phase optimising the deformation of the active contour. For objects undergoing large displacements between frames, we booked a significant improvement of 0.21 in Dice coefficient using our proposed two phase optimisation, compared with regular optimisation.

- Classical external forces attempt to extend the capturing range of an edge by defusing a gradient force field. State-of-the-art forces typically use an isotropic diffusion method. We proposed a novel external force, which uses an anisotropic diffusion process. This allows excluding regions of the image that certainly do not correspond with an object of interest and allows incorporating prior knowledge of the shape of the objects of interest. This new force field was proven useful for the tracking of Arabidopsis thaliana leaves in chlorophyll fluorescence time lapse sequences. After 20 frames the active contours optimised using the proposed force had an average Dice coefficient of 0.92, whereas active contours using other forces only resulted in an average Dice coefficient of 0.74.

- A new and generic external force was proposed. This force is computationally efficient to calculate, i.e. it only requires scanning the images twice in order to propagate edge and force information. The proposed force results

in faster convergence of the active contours, while being computationally much more efficient that other forces: it only requires about a tenth of the calculation time of VFC or GVF.

- We reformulated the active contour as a likelihood maximisation framework instead of an ad hoc energy minimisation framework. Not only is the task of finding appropriate weighting parameters omitted, the proposed tracking framework shows more accurate tracking results as well, even compared with conventional active contours with optimal parameters. For the segmentation of sugar beet seedling leaves an improvement of 0.05 in Dice coefficient was measured. Moreover, all leaves where correctly detected, whereas conventional active contours resulted in the loss of 24 leaves out of 224 leaves.

The work presented in this chapter resulted in a number of international conference publications [De Vylder 09a, De Vylder 09b, De Vylder 10h, De Vylder 10d, De Vylder 10e, De Vylder 11d, De Vylder 13b].

5

Geometric active contours

The previous chapter elaborated on active contours using a parametric representation of the contour, i.e. active contours that represent the contour explicitly using a parametric function: $[0, 1] \mapsto \mathbb{R}^2$. This is an intuitive contour representation, but it is not the only way to represent a contour. A different representation method was proposed in [Malladi 95, Malladi 96, Sethian 99], where a contour is represented using level sets. This approach represents the contour implicitly using a level set function, e.g. $\phi(.) : \mathbb{R}^2 \mapsto \mathbb{R}$. A level set, Θ_c is defined as the set of all points, $\mathbf{x}$, that have the same value for the level set function, i.e.

$$\Theta_c = \{\mathbf{x} \mid \mathbf{x} \in \mathbb{R}^2, \phi(\mathbf{x}) = c\}, \tag{5.0.1}$$

where c is called the level. Although this representation might seem unnecessary complex at first, there are several advantages to this implicit representation:

- The level set function operates on the same domain as the image, so it is possible to calculate a value for each pixel in the image. Therefore, there is no need for contour parameterisation. This is in contrast to the contours used in chapter 4, where the chosen parameterisation will influence the segmentation result.
- This representation does not put any constraint on the topology of the contour. While some attempts have been done for parametric active contours to allow the contours to split in to multiple objects, or to merge two contours,

these attempts are often cumbersome and result in extra rules and decisions that have to be made. These rules have to be checked during each optimisation step, thus resulting in an extra computational burden. In contrast to parametric representations, level sets can change their topology in a natural and trivial way. Fig. 5.1 shows an example of this concept. This example shows a mitotic cell nucleus, i.e. of a cell that is dividing. Such a cell nucleus can be visualised using fluorescence microscopy, see Fig. 5.1.A. This mitotic cell shows two clear parts: the nucleus itself, depicted in light grey, and the clusters of genes, which are shown as two bright regions. Depending on the application, it might be desirable to delineate the full nucleus or only the gene clusters. This can be achieved by the same level set function. Fig 5.1.B shows an example of such a level set function. By changing the level, it is possible to either delineate a single large object (see the second row) or two smaller regions (see the last row).

- The use of intensity information of the complete object has been investigated as a data-fit term for parametric active contours, but optimising the energy term remains a non-trivial task. The main challenge lies in the optimisation of the contour, which is mainly based on information near the edges of the contour, which does not necessarily reflect the full interior/exterior of the contour. This results in slow or bad convergence based on some heuristics [Sundaramoorthi 09, Nakhmani 12]. The geometric representation is more suited to incorporate intensity information since it can optimize the contour indirectly by deforming the level set function. Since the level set function operates on the same domain as the image, intensity information for the full segment can be incorporated in a straightforward manner. We will elaborate on the use of intensity information in the remaining of this chapter

This chapter is arranged as follows: first, we discuss the background and current state-of-the-art of geometric active contours. In section 5.2 we propose a probabilistic geometric active contour framework. We elaborate on the different intensity and shape models that can be used within this Bayesian framework. Next, in section 5.3 we extensively discuss and compare different methods for the optimisation of geometric active contours. Finally, section 5.4 recapitulates and concludes.

5.1 Background

5.1.1 Mumford-Shah piecewise constant active contours

A popular example of intensity based active contour segmentation is based on the Mumford-Shah piecewise constant model [Mumford 89]. This model assumes that

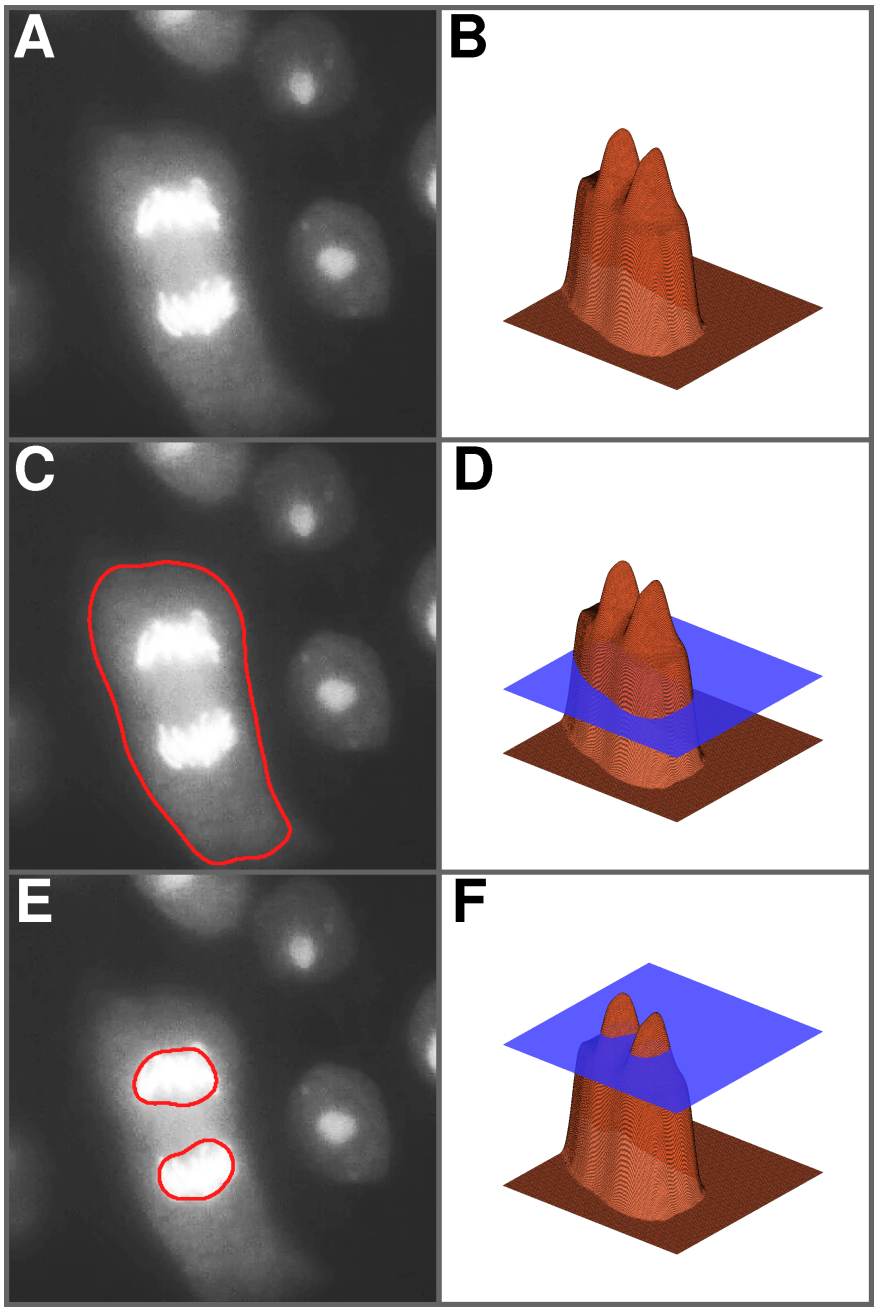


Figure 5.1: An example of the use of level sets to represent a contour. A) a fluorescence microscopic image of a mitotic nucleus. B) a level set function that is used to describe the contours. C) the delineation of the full nucleus, D) the level set representation of the same contour. The contour corresponds to the intersection of the blue plane and the orange level set function. E) the delineation of the genes, F) the level set representation using the same level set function as for (D).

an image is composed of a small number of regions that have a constant intensity. The model itself is independent of the representation of the contour, therefore we define $C(.)$ as a contour representation of interest, Ω_1 as the set of points enclosed by the contour and Ω_2 as all points not enclosed by the contour. The Mumford-Shah model can be used for segmentation by minimising the following energy function:

$$E[C(.)] = \text{length}(C(.)) + \beta \int_{\mathbf{x}\in\Omega_1} (f(\mathbf{x}) - \mu_1)^2 d\mathbf{x} + \beta \int_{\mathbf{x}\in\Omega_2} (f(\mathbf{x}) - \mu_2)^2 d\mathbf{x}, \tag{5.1.1}$$

where $f(.)$ represents the image, β a weighting parameter, while μ_1 and μ_2 correspond respectively to the mean interior and exterior intensity. These values are typically iteratively updated. An optimal contour is calculated based on some estimation of μ_1 and μ_2. After finding an optimal contour, $\mu_1$and μ_2 are updated again. This iterative updating is repeated until convergence is reached. While the Mumford-Shah piecewise constant model is independent of the contour representation, its practical usability is not. In [Sundaramoorthi 09, Nakhmani 12] the contour was represented by a parametric curve, such as in chapter 4. This parametric representation required a heuristic, slowly converging optimising scheme. A more elegant and efficient implementation was proposed in [Chan 99], where the contour is represented implicitly using a level set function. Minimising eq. (5.1.1) then corresponds to finding a level set that minimizes the following energy function:

$$\begin{aligned} E[\phi(.)] &= \text{length}(\phi(\mathbf{x}) = 0) + \beta \int_{\phi(\mathbf{x})\geq 0} (f(\mathbf{x}) - \mu_1)^2 d\mathbf{x} \\ &\quad + \beta \int_{\phi(\mathbf{x})<0} (f(\mathbf{x}) - \mu_2)^2 d\mathbf{x} \\ &= \int_{\mathbf{x}\in\Omega} \left|\nabla H(\phi(\mathbf{x}))\right| d\mathbf{x} + \beta \int_{\mathbf{x}\in\Omega} H(\phi(\mathbf{x}))\,(f(\mathbf{x}) - \mu_1)^2 d\mathbf{x} \\ &\quad + \beta \int_{\mathbf{x}\in\Omega} H(1 - \phi(\mathbf{x}))\,(f(\mathbf{x}) - \mu_2)^2 d\mathbf{x}, \end{aligned} \tag{5.1.2}$$

with Ω representing the image domain, i.e. $\Omega = \Omega_1 \cup \Omega_2$, and where $H(.)$ corresponds to the Heaviside function, i.e.

$$H(x) = \begin{cases} 0 & \text{if } x < 0 \\ 0.5 & \text{if } x = 0 \\ 1 & \text{if } x > 0 \end{cases}. \tag{5.1.3}$$

Note that if $\phi(.)$ is a continuous function, that the $\text{length}(\phi(\mathbf{x}) = 0)$ equals $\int_{\mathbf{x}\in\Omega} \big|\nabla H(\phi(\mathbf{x}))\big| d\mathbf{x}$ [Evans 92]. The discontinuity of the Heaviside hampers the use of Euler Lagrange equations for the optimisation of eq. (5.1.2). In practice this discontinuity is avoided by using a regularised version of H:

$$H_\epsilon(x) = \frac{1}{2}\big(1 + \frac{2}{\pi}\arctan\frac{x}{\epsilon}\big). \tag{5.1.4}$$

As $\epsilon \to 0$ this regularised approximation converges to the Heaviside function. The left of Figure 5.2 shows that the regularised Heaviside function converges towards the actual Heaviside function for a decreasing ϵ. This regularised function is not only continuous, but also differentiable. The right plot of Fig. 5.2 shows the derivative of the regularised Heaviside function for decreasing ϵ. Note that this derivative converges towards a Dirac pulse for $\epsilon \to 0$. Using these regularisations, eq. (5.1.2) can be minimised by treating $\phi(\mathbf{x}, t)$ as a function of time, i.e. the level set function is iteratively optimised using gradient descent, with

$$\frac{\partial\phi(\mathbf{x},t)}{\partial t} = \frac{\partial H_\epsilon(\phi(\mathbf{x},t))}{\partial\phi(\mathbf{x},t)}\Big(\nabla\cdot\frac{\nabla\phi(\mathbf{x},t)}{\big|\nabla\phi(\mathbf{x},t)\big|}) - \beta\,(f(\mathbf{x})-\mu_1)^2 + \beta\,(f(\mathbf{x})-\mu_2)^2\Big). \tag{5.1.5}$$

This results in the following optimality condition:

$$0 = \delta_\epsilon(\phi(\mathbf{x},t))\Big(\nabla\cdot\frac{\nabla\phi(\mathbf{x},t)}{\big|\nabla\phi(\mathbf{x},t)\big|}) - \beta\,(f(\mathbf{x})-\mu_1)^2 + \beta\,(f(\mathbf{x})-\mu_2)^2\Big), \tag{5.1.6}$$

with $\delta_\epsilon(.)$ a regularised version of the one-dimensional Dirac function. Note that for $\epsilon \to 0$ the Dirac function dominates this optimality condition. This in combination with the fact that the derivative of a Dirac function is not positive for its full domain, proves that eq. (5.1.2) is a non-convex function. As such the energy function proposed in eq. (5.1.2) is hampered by local optima. As a consequence, the contour found by solving eq. (5.1.6) depends on the initialisation of the level set function.

5.1.2 Globally optimal active contours

In [Chan 06] a segmentation method similar to the method discussed in the previous subsection was proposed. However, instead of using level sets, a different non-parametric representation is used: a characteristic function. A characteristic function, $u(\mathbf{x})$, has a zero value if the pixel does not belong to the segment, otherwise it results in 1. The segmentation result of [Chan 06] is calculated by minimising the following energy function:

$$E\big[u(.)\big] = \big|\nabla u(.)\big| + \beta\big\langle u(.), r(.)\big\rangle, \tag{5.1.7}$$

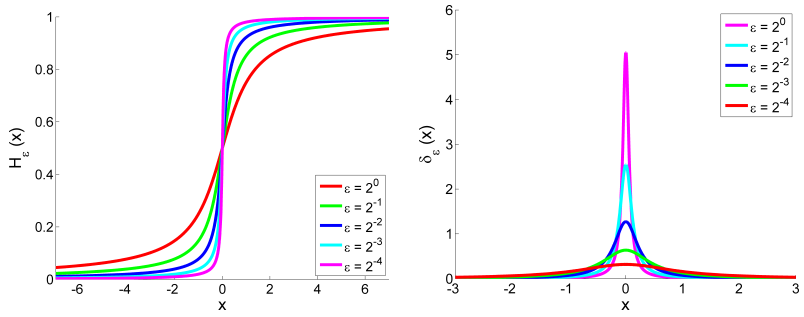


Figure 5.2: The regularised Heaviside function for different ϵ. Left, the actual regularised Heaviside function, on the right the derivative of the regularised Heaviside function.

with β is a weighting parameter used to tune the influence of the data-fit term in relation to the total variation regularisation and

$$r(\mathbf{x}) = (\mu_1 - f(\mathbf{x}))^2 - (\mu_2 - f(\mathbf{x}))^2, \tag{5.1.8}$$

where $\langle ., . \rangle$ is the inner product and μ_1, μ_2 are respectively the expected foreground and background intensity.. Note that if μ_1 and μ_2 are constant, e.g. calculated from a training data set, then eq. (5.1.7) represents a convex function. In order to allow the use of efficient, global convex optimisers, the set of possible solution, i.e. $u(.)$, has to be a convex set. Hence Chan et al. [Chan 06] relaxed the constraint that $u(.)$ is a characteristic function, instead using the set, S, of functions with codomain equal to the convex region $[0, 1]$. This results in the following convex optimisation problem:

$$\hat{u}(.) = \underset{u(.) \in S}{\arg\min} \left|\nabla u(.)\right| + \beta \langle u(.), r(.) \rangle. \tag{5.1.9}$$

Furthermore, this formulation relates to the popular and widely used Mumford-Shah piecewise constant model, eq. (5.1.1). This can be seen by comparing the optimality condition in eq. (5.1.6) with the steady state of the gradient flow corresponding to this convex energy function, i.e. the optimality condition of eq. (5.1.7) which is equal to:

$$0 = \left(\nabla \cdot \frac{\nabla u(\mathbf{x}, t)}{\left|\nabla u(\mathbf{x}, t)\right|}\right) - \beta \left(f(\mathbf{x}) - \mu_1\right)^2 + \beta \left(f(\mathbf{x}) - \mu_2\right)^2. \tag{5.1.10}$$

By replacing $u(.)$ with $u(.) - 0.5$, we convert the characteristic function to a level set function, where the zero level corresponds to the same contour represented by $u(.)$. By multiplying $\delta_\epsilon(u(.) - 0.5)$ with the left and right part of eq. (5.1.10), we get the steady state corresponding with the convex energy of eq. (5.1.9). In other words, an optimum of this convex energy coincides with the steady state of

the gradient flow of the original Mumford-Shah energy function (section 5.1.1). Note that by relaxing the constraint that the function that minimizes eq. (5.1.9) is a characteristic function, that $\hat{u}(\mathbf{x})$ can have any value between 0 and 1. As a consequence, the resulting function not necessarily has to represent a crisp segmentation. A binary segmentation result can be obtained by thresholding $\hat{u}(\mathbf{x})$, i.e.

$$\text{thresh}\left(\hat{u}(\mathbf{x});\alpha\right) = \begin{cases} 1 & \text{if } \hat{u}(\mathbf{x}) > \alpha \\ 0 & \text{otherwise} \end{cases} \tag{5.1.11}$$

for $\alpha \in [0,1]$. Although different values for α can result in different segmentation results, it can be proven that irrespective of α, $\text{thresh}\left(\hat{u}(\mathbf{x});\alpha\right)$ itself is also a global minimiser for the energy in eq. (5.1.7) and by extension for the energy function of the Mumford-Shah model [Bresson 07a, Pock 10]. It is possible that a different α results in a different binary segmentation, however the energy of these different binary segmentations will be equal. So there exists a binary segmentation that minimises eq. (5.1.7), however it does not have to be unique.

5.2 Probabilistic geometric active contours

While the energy function of eq. (5.1.9) has been proven useful for segmentation [Sethian 99,Sethian 02,Gong 05,Papandreou 07,Mille 09], it is non-straightforward to define optimal weighting parameters. Furthermore many alternative regularisation terms and data-fit terms have been proposed. This excess on tuning possibilities makes it difficult to define an optimal energy function, $f(.)$, for a given segmentation problem. In order to avoid ad hoc energy terms and weights, a more statistical founded approach is desirable. Geometric active contours with a probabilistic interpretation have been reported in literature [Cremers 08, Chen 10, YanFeng 11, Stuhmer 13]. However, these methods either limit the usability by imposing strong but restrictive shape priors [Cremers 08, Stuhmer 13], or start from a probabilistic interpretation for the data-fit but use generic shape energies without any probabilistic interpretation, thus returning to ad hoc parameter tuning [Chen 10]. In this section we propose a generic statistical framework for active contour segmentation that circumvents ad hoc parameter tweaking. This Bayesian framework approaches segmentation as an inverse problem. We model an image as the output of a function, $g(.)$, which maps pixels belonging to an object of interest to a specific intensity, while mapping background pixels to a different intensity. This model of course is not perfect, e.g. there can be noise in the image, or blur, or irrelevant objects might clutter the image, etc. All these discrepancies on the model can be added as noise, $n(.)$. Then the image $f(.)$ can be defined as:

$$f(\mathbf{x}) = g(u(\mathbf{x})) + n(u(\mathbf{x}), \mathbf{x}). \tag{5.2.1}$$

In the remainder of this section we will refer to the combination of the noise model and the mapping function as the image model. Using this image model, segmentation results in finding the characteristic function $u(.)$, that is most likely to result in the given image $f(.)$. This can be achieved by using Bayesian estimators. In this section we propose to define the segmentation result as the Maximum A Posteriori estimator (MAP), i.e.

$$\hat{u}(\mathbf{x}) = \underset{u(.)\in S}{\arg\max} \log \mathrm{p}_{\mathrm{U}}\big(u(\mathbf{x})\big) + \log \mathrm{p}_{\mathrm{F|U}}\big(f(\mathbf{x}) \mid u(\mathbf{x})\big). \tag{5.2.2}$$

For readability we write the random variable as U instead of U($\mathbf{x}$), in situations where we want to emphasize that the random variable is a function and not a value of a function we will explicitly write U(.). Note that $\mathrm{p}_{\mathrm{F|U}}\big(f(\mathbf{x}) \mid u(\mathbf{x})\big)$ depends on the image model. The quality of the segmentation resulting from this MAP estimator strongly depends both on the likelihood of the image given $u(.)$, and on the prior knowledge of $u(.)$. We will discuss different priors and image models in the next two subsections.

5.2.1 The image model

In this subsection we will focus on the image model and its impact on the likelihood of the image given the segmentation, i.e. $\mathrm{p}\big(f(.) \mid u(.)\big)$ in eq. (5.2.2). We will describe several models that can be used for segmentation. For the sake of simplicity we will assume that the segment and the background have the same type of image model. However this is just for didactic purpose, in practice foreground and background can use different image models according to the intrinsic properties of the images.

5.2.1.1 Piecewise constant Mumford-Shah

In section 5.1 the piecewise constant Mumford-Shah model was already presented. This model assumes a simple data-fit model, i.e. that an image has a constant value for each segment pixel, μ_1, and a different constant value for each background pixel , μ_2. This results in the following image model:

$$g(a) = \mu_1\, a + \mu_2\, (1 - a), \tag{5.2.3}$$

with $a \in \{0, 1\}$. As mentioned before, we assume that the noise statistics can differ between segment pixels and background. If we model the noise for both segment and background by a Normal distribution, then the likelihood of $f(.)$ corresponds to:

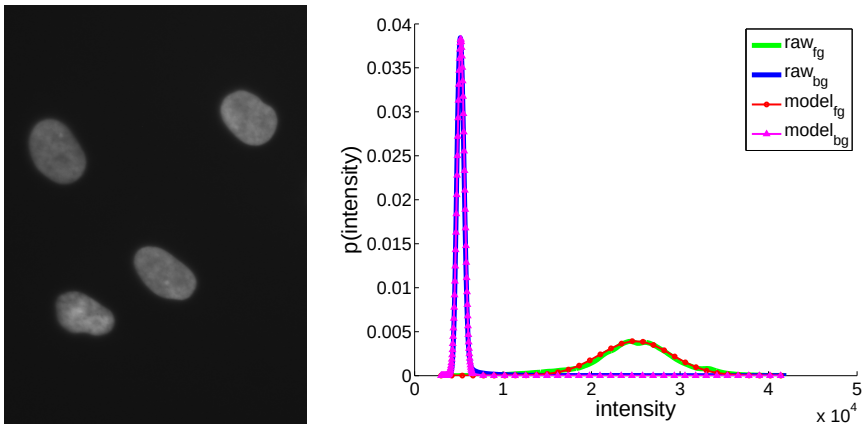


Figure 5.3: An example of biological images where the Mumford-Shah model holds. On the left cell nuclei imaged using fluorescence microscopy. On the right, the likelihood of the intensity of a pixel belonging to the nucleus and for pixels belonging to the background.

$$\mathrm{p}_{\mathrm{F|U}}\big(f(\mathbf{x}) \mid u(\mathbf{x})\big) = \begin{cases} \frac{1}{\sqrt{2\pi\sigma_{i1}^2}} e^{\frac{-(f(\mathbf{x})-\mu_1)^2}{2\sigma_{i1}^2}} & \text{if } u(\mathbf{x}) = 1 \\ \frac{1}{\sqrt{2\pi\sigma_{i2}^2}} e^{\frac{-(f(\mathbf{x})-\mu_2)^2}{2\sigma_{i2}^2}} & \text{if } u(\mathbf{x}) = 0 \end{cases}, \tag{5.2.4}$$

where σ_{i1}^2 and σ_{i2}^2 correspond to the standard deviation of the noise on the intensity of respectively the segment and background. Both the piecewise constant assumption and the Gaussian noise assumption result in a rather simplistic model. Despite its simplicity, it proves to be a valid and useful model for numerous applications. Fig. 5.3 shows an example of such an application. On the left cell nuclei are visualised using fluorescence microscopy. The right part of the figure shows the likelihood of a certain intensity for both cell nuclei and for the background of the left image. These likelihoods are quite accurately modelled using a Gaussian model, as can be seen by the cyan and red curves.

Since we assume that $u(.)$ represents a characteristic function, i.e. $u(\mathbf{x}) \in \{0, 1\}$, we can rewrite the log likelihood as follows:

$$\begin{aligned}\log \mathrm{p}_{\mathrm{F|U}}\big(f(\mathbf{x}) \mid u(\mathbf{x})\big) &= u(\mathbf{x})\left(\frac{-(f(\mathbf{x})-\mu_1)^2}{2\sigma_{i1}^2} + \log\frac{1}{\sqrt{2\pi\sigma_{i1}^2}}\right) \\ &\quad +(1-u(\mathbf{x}))\left(\frac{-(f(\mathbf{x})-\mu_2)^2}{2\sigma_{i2}^2} + \log\frac{1}{\sqrt{2\pi\sigma_{i2}^2}}\right) \\ &= u(\mathbf{x})\left(\frac{1}{2\sigma_{i2}^2}(f(\mathbf{x})-\mu_2)^2 - \frac{1}{2\sigma_{i1}^2}(f(\mathbf{x})-\mu_1)^2\right) \\ &\quad +u(\mathbf{x})\left(\log\frac{1}{\sqrt{2\pi\sigma_{i1}^2}} - \log\frac{1}{\sqrt{2\pi\sigma_{i2}^2}}\right) + c,\end{aligned} \tag{5.2.5}$$

with c a constant independent of $u(\mathbf{x})$, thus we can ignore this constant with regard to the MAP estimator. Note that for $\sigma_1 = \sigma_2$ this equates to the same data fit term as used in section 5.1.2. The benefit of this approach is that it gives a statistical interpretation for the weighting parameter β in eq. (5.1.9). This probabilistic interpretation however, is more general than the model used in section 5.1.2, i.e. the model allows different variances for segment and background noise. This is important to allow accurate modelling, which eventually results in more accurate segmentation. This is illustrated with an example in Fig. 5.4. Figure 5.4.A shows a microscopic image containing pathological cell nuclei. These pathologies result in aberrant shapes and non-crisp cell boundaries. Fig. 5.4.B shows the segmentation result using the classical Mumford Shah model, i.e. the method discussed in section 5.1.2, where the variance of both foreground and segment noise are considered equal. This clearly results in erroneous segmentation of several nuclei (delineated in red). By relaxing the constraint that the noise variance is equal for the segments and the background, the cell nuclei are accurately segmented (delineated in green in Fig. 5.4.C).

If this image model is used in a maximum likelihood estimator (ML) instead of a MAP estimator, i.e. without any prior knowledge on $u(.)$, the optimisation of eq. (5.2.5) can efficiently be implemented using hard thresholding.

5.2.1.2 General stationary image models

The piecewise constant Mumford-Shah model has the benefit that it is easy to calculate the parameters for the log likelihood, i.e. estimating only four parameters can be achieved with relatively little training data. The downside is that the model is sometimes too restrictive for certain applications. However, the approach used in the previous subsection can be generalised into the following image model

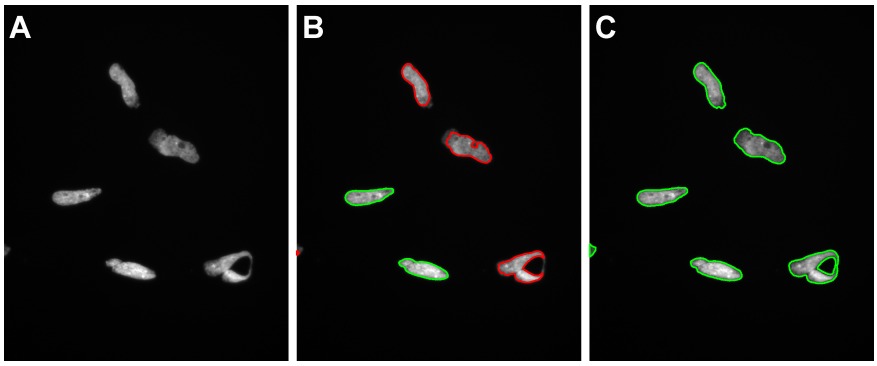


Figure 5.4: An example of segmentation based on the Mumford-Shah assumption. A) a fluorescent microscopic image showing pathological cell nuclei, B) the segmentation result assuming that the segment and the background have the same noise variance [Chan 06], C) the segmentation result with the proposed model with correctly modelled variances [De Vylder 14c].

$$\log \mathrm{p}_{\mathrm{F}|\mathrm{U}}\big(f(\mathbf{x}) \mid u(\mathbf{x})\big) = u(\mathbf{x})\big(\log \mathrm{p}_{\mathrm{F}|\mathrm{H}_1}(f(\mathbf{x}) \mid H_1) - \log \mathrm{p}_{\mathrm{F}|\mathrm{H}_2}(f(\mathbf{x}) \mid H_2)\big), \tag{5.2.6}$$

where H_1 represents the hypothesis that $\mathbf{x}$ is a background pixel and H_2 corresponds to the hypothesis that it represents a segment pixel. Within this generalised stationary model the probability density functions are not limited to normal distributions. The probability density functions are typically estimated from a training dataset. This general model is also well suited for multi-channel images [Nieuwenhuis 13a]. Fig. 5.5 shows an example of such a multi-channel microscopic image. This image shows the root of an Arabidopsis plant. In Fig 5.5.A the cell walls are marked using a specific fluorescent protein, whereas in Fig. 5.5.B the cell nuclei are visualised. Normally this image should only show bright blobs, since the fluorescent protein is only present within the nuclei. However, the image is corrupted with bleed trough, i.e. some light emitted by the fluorescent dyes used to mark the cell walls is not filtered out. As a result the cell walls become partially visible within this channel. Fig. 5.5.C shows the likelihood that a cell nucleus is present at position $\mathbf{x}$, i.e. $\log \mathrm{p}_{\mathrm{F}|\mathrm{H}_1}(f(\mathbf{x}) \mid H_1)$. Fortunately this likelihood is unaffected by the bleed through, enabling more accurate segmentation.

Even if an extensive annotated dataset is not available for the training of the probability density functions, this general model can still be used. In such a situation the probability density functions have to be estimated based on the image. To model the likelihoods for both hypothesis H_1 and H_2, we need to make a distinction between intensities belonging to foreground and background. To achieve this, we propose to approximate the image histogram using a Gaussian mixture model,

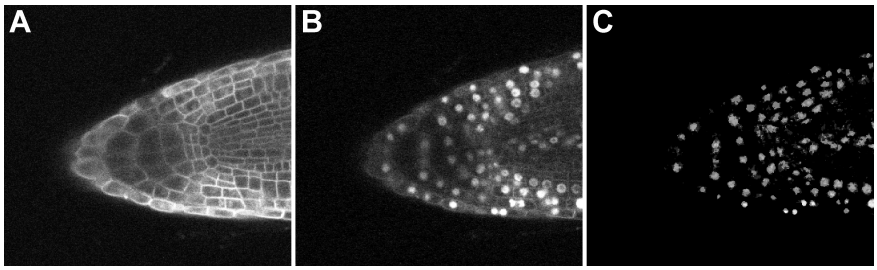


Figure 5.5: An example of a multi-channel microscopy image of the root of an Arabidopsis plant. A) the cell walls marked using a fluorescent dye, B) The cell nuclei marked using a different fluorescent dye (DAPI), C) the likelihood image for cell nuclei [De Vylder 10g].

i.e.

$$h(\mathbf{i}) \approx \sum_{j=1}^{l} \alpha_j \mathcal{N}(\mathbf{i}; \boldsymbol{\mu}_j, K_j), \tag{5.2.7}$$

where α_j is a weighting parameter such that $\alpha_j \in [0, 1]$ and $\sum_{j=1}^{l} \alpha_j = 1$ and where $\mathcal{N}(.; \boldsymbol{\mu}_j, K_j)$ represents the probability density function of a Gaussian distribution with mean equal to $\boldsymbol{\mu}_j$ and K_j the covariance matrix. A human operator can now make a selection of Gaussian distributions that correspond to foreground pixels. The mixture of these Gaussians corresponds to $\mathrm{p}_{\mathrm{F|H_2}}(f(\mathbf{x}) \mid H_2)$, whereas the mixture of all other Gaussian components corresponds to $\mathrm{p}_{\mathrm{F|H_1}}(f(\mathbf{x}) \mid H_1)$. The number of Gaussian components used, l in eq. (5.2.7), could be defined by a human operator or could be iteratively incremented until the approximation of the Gaussian mixture model resembles the image histogram sufficiently. The probability parameters $\boldsymbol{\mu}_j$ and K_j can be calculated using the expectation maximisation algorithm [Shang 11].

As an example we will segment bacteria in a transmission electron microscopy image. Fig. 5.6.A shows the segmentation result of the original convex Chan-Vese active contour [Chan 06](discussed in section 5.1.2). As can be seen, the conventional convex active contour assuming two normal distributions results in erroneous segmentation since the light grey interior of the bacteria is discarded by the method. The result of our proposed method can be seen in Fig. 5.6.B. To get this result, the image histogram was approximated by a mixture of 3 Gaussian components. The Gaussian distribution with highest mean was assumed to correspond to the probability distribution of the background. The two other Gaussian distributions are considered to represent two modalities of foreground pixels. The proposed method results in correct segmentation of the complete bacteria and this without a training step for the required probability distributions.

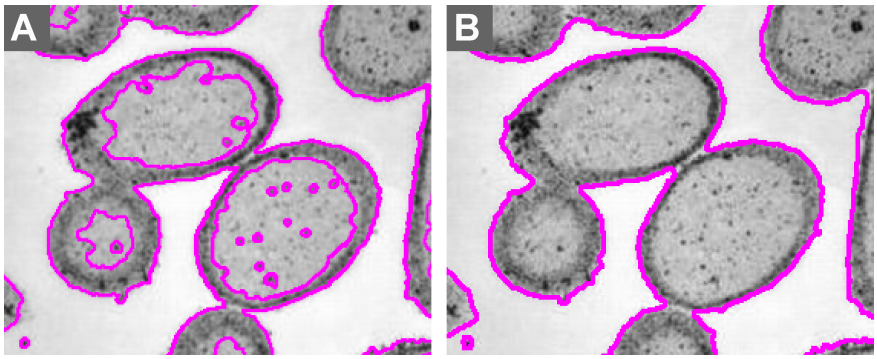


Figure 5.6: An example of segmenting bacteria in electron microscopy images. A) the segmentation result using two normal distributions, B) segmentation using the generalised data-fit term based on a mixture of Gaussians model [De Vylder 11g].

5.2.1.3 Non-stationary image models

The previous two image models assumed that all pixel intensities are independent of their location, which is a harsh constraint that is violated by many applications. Therefore a more powerful image model can be considered, e.g. where the image model does not impose a constant value, but where the segment and background intensities correspond to a specific function:

$$g(\mathbf{x}) = \begin{cases} g_1(\mathbf{x}) & \text{if } u(\mathbf{x}) = 1 \\ g_2(\mathbf{x}) & \text{if } u(\mathbf{x}) = 0 \end{cases}. \tag{5.2.8}$$

Template based image models In the situation where $g_1(.)$ and $g_2(.)$ are known, this results in the following likelihood (assuming Gaussian white noise):

$$\mathrm{p}_{\mathrm{F|U,G_1,G_2}}\big(f(\mathbf{x}) \mid u(\mathbf{x}),\, g_1(\mathbf{x}), g_2(\mathbf{x})\big) = \begin{cases} \frac{1}{\sqrt{2\pi\sigma_{i1}^2}} e^{\frac{-(f(\mathbf{x})-g_1(\mathbf{x}))^2}{2\sigma_{i1}^2}} & \text{if } u(\mathbf{x}) = 1 \\ \frac{1}{\sqrt{2\pi\sigma_{i2}^2}} e^{\frac{-(f(\mathbf{x})-g_2(\mathbf{x}))^2}{2\sigma_{i2}^2}} & \text{if } u(\mathbf{x}) = 0 \end{cases}. \tag{5.2.9}$$

Fig. 5.7 shows an example where such an image model can be used. In these images the visual symptoms of tobacco mosaic virus should be detected. This infection manifests as bright blobs on the leaf [Chaerle 00, Chaerle 01], as can be seen in Fig. 5.7.B. Only using colour or brightness would result in inaccurate segmentation, since the nerves of the leaf have a similar colour and brightness. However, by using a template based image model this problem becomes much

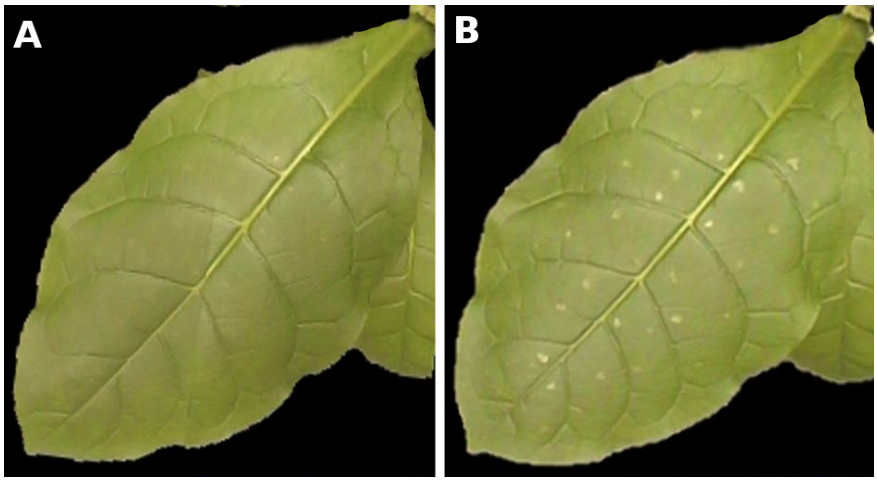


Figure 5.7: An example of detecting visual symptoms tobacco mosaic virus on leafs. A) a leaf before infection becomes visible. This image can serve as a background model. B) the same leaf when the infection is visible. Images by courtesy of Dominique Van Der Straeten, Laboratory of Functional Plant Biology, Faculty of Sciences, Ghent University.

easier. In this application the template can be defined by imaging the leaf before infection, see fig. 5.7.A.

Note that non-stationary models are not restricted to foreground-background segmentation. Fig. 5.8.A shows an example of a different application where non-stationary models can be used. In this example the cell walls of an Arabidopsis root are marked using a fluorescent dye. Note that the intensity of the cell walls depends on the location within the root, i.e. outer cells have a higher intensity than the cells in the centre of the root. This is mainly due to their 3D location in the root. The inner cells are less in focus than the outer cells. Using simple morphological operators local maxima and minima can be detected. These local extrema are considered as "certain" cell wall or background pixels, depending if they correspond to a maximum or a minimum. Using polyharmonic spline interpolation, the intensity of these "certain pixels" can be used to estimate continuous background and foreground functions, i.e. respectively $g_2(.)$ and $g_1(.)$ [Mat 14, De Vylder 10g]. Figures 5.8.B-C show an example of such a background and foreground function. Fig. 5.8.D shows the log-likelihood of a pixel belonging to the foreground minus the log-likelihood that a pixel belongs to the background, i.e. $\log \mathrm{p}_{\mathrm{F}|\mathrm{H}_1,\mathrm{G}_1}\big(f(\mathbf{x}) \mid H_1, g_1(\mathbf{x})\big) - \log \mathrm{p}_{\mathrm{F}|\mathrm{H}_2,\mathrm{G}_2}\big(f(\mathbf{x}) \mid H_2, g_2(\mathbf{x})\big)$ under the Gaussian assumption of eq. (5.2.9) and with $H_1(\mathbf{x})$ the hypothesis that the pixel $\mathbf{x}$ corresponds to the segment, whereas $H_2(\mathbf{x})$ is the hypothesis that the pixel belongs to the background. Note that while the intensity of cell walls is not constant in the original fluorescent image, this is much more the case in Fig. 5.8.D,

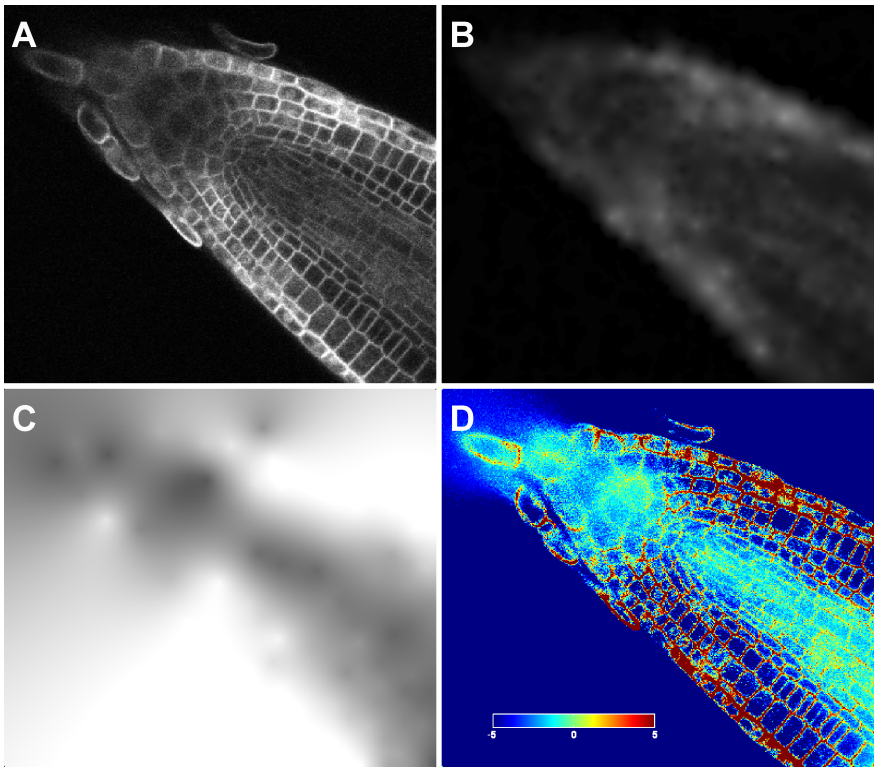


Figure 5.8: An example of a microscopic image of a root, with the cell walls are marked with a fluorescent protein. A) the original fluorescent microscopic image, B) the estimate background function, C) the estimated foreground function, D) the likelihood map of pixels corresponding to the foreground minus the likelihood that a pixel belongs to the background.

e.g. cell walls not at the border of the plant root have a similar value as the cell walls at the boundary. The central part of the root shows high values at the cell walls but values near zero at the centre of cells. This is not unexpected since the central regions have a lower signal to noise ratio. Fortunately the segmentation result does not only rely on the likelihoods of the intensity, but also incorporates prior knowledge on the expected shape of the segments. These priors will typically correct erroneous segmented pixels that have a similar likelihood for background and segment. For a detailed discussion on shape priors we refer to section 5.2.2.

Polynomial image model In some cases the function $g_1(.)$ and $g_2(.)$ are known, e.g. a fixed background, however in many cases these functions are object depend. An example of such a situation is shown in Fig. 5.9. Here the intensity of the leaf in the chlorophyll fluorescence image depends on the orientation and depth of the

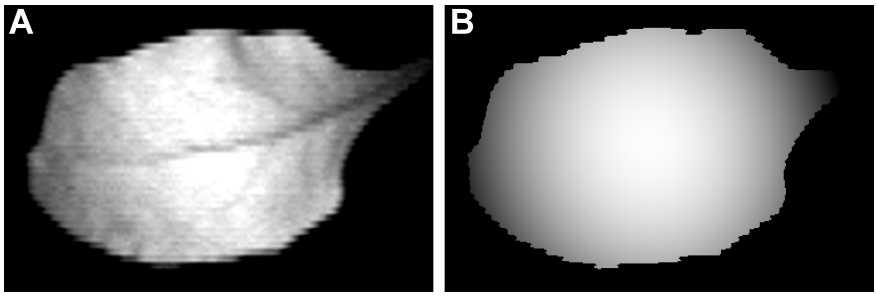


Figure 5.9: An example of a polynomial image model. A) a chlorophyll fluorescence image of an isolated Arabidopsis leaf. B) a polynomial of second degree fitted to the intensity of the leaf.

leaf. This is of course non-stationary, thus the piecewise constant model of section 5.2.1.1 cannot be used. This problem can be tackled by estimating the functions $g_1(.)$ and $g_2(.)$ concurrently with $u(.)$. To overcome overfitting, some prior knowledge on these function should be incorporated. Restricting $g_i(.)$ to a low degree polynomial function has been proven a useful constraint for segmentation [Deboeverie 13a, Deboeverie 13b]. Fig. 5.9.B shows an example of a two dimensional polynomial of second degree fitted to the leaf. This polynomial model is less hampered by the global intensity fluctuation of the leaf than a constant intensity model. Furthermore, prior knowledge on the parameters of the polynomial can be added in a straightforward way. We will explain this with a simple example where the foreground is modelled using a polynomial of order 2, i.e.

$$g_1(x, y) = a_1 + a_2 x + a_3 y + a_4 x^2 + a_5 xy + a_6 y^2. \tag{5.2.10}$$

For brevity, we refer to the parameters of this polynomial as a vector, $\mathbf{a_{p1}} = [a_1, a_2, ..., a_6]$. Note that $g_1(.)$ is fully defined by the parameter $\mathbf{a_{p1}}$. So in order to model the likelihood that a certain foreground function $g_1(.)$ occurs, it suffices to model its parameters. As an example we assume that the parameters of the foreground polynomial have a multivariate normal likelihood, this results in the following likelihood for $g_1(.)$:

$$\mathrm{p}_{\mathrm{G}_1(.)}\big(g_1(.)\big) = (2\pi)^{-\frac{k}{2}} \left|\Sigma_{p1}\right|^{-\frac{1}{2}} e^{-\frac{1}{2}(\mathbf{a_{p1}}-\boldsymbol{\mu_{p1}})' \Sigma_{p1}^{-1} (\mathbf{a_{p1}}-\boldsymbol{\mu_{p1}})}, \tag{5.2.11}$$

with $\boldsymbol{\mu_{p1}}$ the expected values for the parameters $\mathbf{a_{p1}}$ and Σ_{p1} the covariance matrix. The expected values and covariance matrix could either be estimated from a training dataset, or in the case of video analysis based on the parameters used by previous frames. Using this prior knowledge on $g_1(.)$ (and a similar model for the background) results in the following log likelihood for the MAP estimator:

$$
\begin{aligned}
& \log \mathrm{p}_{\mathrm{G}_1(.)}(g_1(.))\ \mathrm{p}_{\mathrm{G}_2(.)}(g_2(.))\ \mathrm{p}_{\mathrm{F|U,G_1,G_2}}\big(f(\mathbf{x}) \mid u(\mathbf{x}), g_1(\mathbf{x}), g_2(\mathbf{x})\big) \\
= \quad & u(\mathbf{x})\left(\frac{-(f(\mathbf{x})-g_1(\mathbf{x}))^2}{2\sigma_{i1}^2} + \log\frac{1}{\sqrt{2\pi\sigma_{i1}^2}}\right) \\
& +(1-u(\mathbf{x}))\left(\frac{-(f(\mathbf{x})-g_2(\mathbf{x}))^2}{2\sigma_{i2}^2} + \log\frac{1}{\sqrt{2\pi\sigma_{i2}^2}}\right) \\
& -\frac{1}{2}(\mathbf{a_{p1}}-\boldsymbol{\mu}_{\mathbf{p1}})'\Sigma_{p1}^{-1}(\mathbf{a_{p1}}-\boldsymbol{\mu}_{\mathbf{p1}}) - \frac{1}{2}(\mathbf{a_{p2}}-\boldsymbol{\mu}_{\mathbf{p2}})'\Sigma_{p2}^{-1}(\mathbf{a_{p2}}-\boldsymbol{\mu}_{\mathbf{p2}}) \\
& +\log(2\pi)^{-\frac{k}{2}}\left|\Sigma_{p1}\right|^{-\frac{1}{2}} + \log(2\pi)^{-\frac{k}{2}}\left|\Sigma_{p2}\right|^{-\frac{1}{2}} \\
= \quad & u(\mathbf{x})\left(\frac{1}{2\sigma_{i2}^2}(f(\mathbf{x})-g_2(\mathbf{x}))^2 - \frac{1}{2\sigma_{i1}^2}(f(\mathbf{x})-g_1(\mathbf{x}))^2\right) \\
& -\frac{1}{2}(\mathbf{a_{p1}}-\boldsymbol{\mu}_{\mathbf{p1}})'\Sigma_{p1}^{-1}(\mathbf{a_{p1}}-\boldsymbol{\mu}_{\mathbf{p1}}) - \frac{1}{2}(\mathbf{a_{p2}}-\boldsymbol{\mu}_{\mathbf{p2}})'\Sigma_{p2}^{-1}(\mathbf{a_{p2}}-\boldsymbol{\mu}_{\mathbf{p2}}) \\
& +u(\mathbf{x})\left(\log\frac{1}{\sqrt{2\pi\sigma_{i1}^2}} - \log\frac{1}{\sqrt{2\pi\sigma_{i2}^2}}\right) + \frac{1}{2\sigma_{i2}^2}(f(\mathbf{x})-g_2(\mathbf{x}))^2 + c,
\end{aligned}
\tag{5.2.12}
$$

where c is a constant independent of the polynomial functions $g_i(.)$ and of $u(.)$. Note that this is no longer convex. The MAP estimator of this problem can be calculated by iteratively alternating between estimating $\hat{u}(.)$ for fixed $\hat{g}_1(.)$ and $\hat{g}_2(.)$, then fixing $\hat{u}(.)$ and estimating $\hat{g}_1(.)$ and $\hat{g}_2(.)$ [Bresson 09, Setzer 09]. This process is repeated until convergence is reached.

The Mumford-Shah piecewise smooth model The assumption that the $g_i(.)$ functions correspond to low degree polynomials results in a powerful image model, however the constraint is still too limiting for some applications, e.g. a full plant rosette typically cannot be modelled by a single low degree polynomial. Therefore the piecewise smooth Mumford-Shah model allows an broader set of functions: smooth, i.e. functions with no abrupt changes. This smoothness prior can be expressed using the magnitude of the gradient. The smoothness of $g_i(.)$ can be modelled using many different probability distributions, e.g. Laplacian, Gaussian, etc. By modelling the smoothness with a normal distribution we retrieve the original Mumford Shah piecewise smooth model [Mumford 89]. More specific, with variance σ_{s1}^2, this results in the following prior:

$$
\mathrm{p}_{\mathrm{G}_1}(g_1(\mathbf{x})) \quad = \quad \frac{1}{\sqrt{2\pi\sigma_{s1}^2}} e^{\frac{-|\nabla g_1(\mathbf{x})|^2}{2\sigma_{s1}^2}}. \tag{5.2.13}
$$

This smoothness prior can then be incorporated in the MAP estimator, resulting in the following likelihood:

$$
\begin{aligned}
& \log \mathrm{p}_{\mathrm{G}_1}(g_1(\mathbf{x}))\ \mathrm{p}_{\mathrm{G}_2}(g_2(\mathbf{x}))\ \mathrm{p}_{\mathrm{F|U,G_1,G_2}}\big(f(\mathbf{x}) \mid u(\mathbf{x}), g_1(\mathbf{x}), g_2(\mathbf{x})\big) \\
= \quad & u(\mathbf{x})\left(\frac{-(f(\mathbf{x})-g_1(\mathbf{x}))^2}{2\sigma_{i1}^2} + \log\frac{1}{2\pi\sigma_{i1}\sigma_{s1}}\right) \\
& +(1-u(\mathbf{x}))\left(\frac{-(f(\mathbf{x})-g_2(\mathbf{x}))^2}{2\sigma_{i2}^2} + \log\frac{1}{2\pi\sigma_{i2}\sigma_{s2}}\right) \\
& -\frac{|\nabla g_1(\mathbf{x})|^2}{2\sigma_{s1}^2} - \frac{|\nabla g_2(\mathbf{x})|^2}{2\sigma_{s2}^2} \\
= \quad & u(\mathbf{x})\left(\frac{1}{2\sigma_{i2}^2}(f(\mathbf{x})-g_2(\mathbf{x}))^2 - \frac{1}{2\sigma_{i1}^2}(f(\mathbf{x})-g_1(\mathbf{x}))^2\right) \\
& +u(\mathbf{x})\left(\log\frac{1}{2\pi\sigma_{i1}\sigma_{s1}} - \log\frac{1}{2\pi\sigma_{i2}\sigma_{s2}}\right) \\
& +\left(\frac{|\nabla g_2(\mathbf{x})|^2}{2\sigma_{s2}^2} - \frac{|\nabla g_1(\mathbf{x})|^2}{2\sigma_{s1}^2}\right) + \frac{1}{2\sigma_{i2}^2}(f(\mathbf{x})-g_2(\mathbf{x}))^2 + c, \quad (5.2.14)
\end{aligned}
$$

with c a constant independent of the smooth functions $g_i(.)$ and the characteristic function $u(.)$.

Texture based image models While the polynomial model of eq. (5.2.12) and the smooth model of eq. (5.2.14) are well suited for segments that gradually change in intensity, the model in eq. (5.2.9) is much more generic. The functions $g_i(.)$ are not restricted to smooth functions, but can also correspond to the texture of a segment. Several sparse representations have been proposed that model local intensity variation, i.e. wavelets, steerable pyramids, shearlets, wave atoms, etc. [Van de Wouwer 99, Rooms 05b, Demanet 07, Goossens 10, Goossens 13]. These sparse transforms, S, typically correspond to an algebraicframe where a synthetic texture image corresponds to a linear combination of the frame, i.e. $g_i(.) = S\,\mathbf{t_i}$, for a given $\mathbf{t_i}$. To prevent over-fitting, i.e. to prevent setting $\mathbf{t_i}$ such that $g_i(.)$ perfectly corresponds to the segment, noise inclusive, instead of only the texture of the segment, an extra prior on $\mathbf{t_i}$ should be incorporated. For many applications the likelihood of the components of $\mathbf{t}$ is accurately modelled using a generalised Laplacian distribution [Goossens 10, Aelterman 14]:

$$\mathrm{p}_{\mathrm{G_i}(.)}(g_i(.)) = \prod_j \frac{\beta_{ij}}{2\alpha_{ij}\Gamma(\frac{1}{\beta_{ij}})} e^{-\left(\frac{|\mathbf{t_i}[j]|}{\alpha_{ij}}\right)^{\beta_{ij}}}, \tag{5.2.15}$$

where j depicts the different functions of the frame S, while $\Gamma(x)$ represents the Gamma function, i.e. $\Gamma(x) = \int_0^\infty s^{x-1}e^{-s}ds$. Note that the parameters of the generalised Laplacian can be different for each basis function of the frame. This allows making textures with features at a specific scale or direction more or less likely. Such a texture generation results in the following log likelihood:

$$\begin{aligned}
& \mathrm{p}_{\mathrm{G}_1(.)}(g_1(.))\, \mathrm{p}_{\mathrm{G}_2(.)}(g_2(.))\ \log \mathrm{p}_{\mathrm{F|U,G_1,G_2}}\big(f(\mathbf{x}) \mid u(\mathbf{x}), g_1(\mathbf{x}), g_2(\mathbf{x})\big) \\
= \;& u(\mathbf{x})\left(\frac{1}{2\sigma_{t2}^2}(f(\mathbf{x}) - g_2(\mathbf{x}))^2 - \frac{1}{2\sigma_{t1}^2}(f(\mathbf{x}) - g_1(\mathbf{x}))^2\right) \\
& + \sum_j \left(\left(\frac{|\mathbf{t_2}[j]|}{\alpha_{2j}}\right)^{\beta_{2j}} - \left(\frac{|\mathbf{t_1}[j]|}{\alpha_{1j}}\right)^{\beta_{1j}}\right) \\
& + u(\mathbf{x})\left(\log\frac{1}{\sqrt{2\pi\sigma_{t1}^2}} - \log\frac{1}{\sqrt{2\pi\sigma_{t2}^2}}\right) + c \\
= \;& u(\mathbf{x})\left(\frac{1}{2\sigma_{t2}^2}(f(\mathbf{x}) - S\,\mathbf{t_2})^2 - \frac{1}{2\sigma_{t1}^2}(f(\mathbf{x}) - S\,\mathbf{t_1})^2\right) \\
& + \sum_j \left(\left(\frac{|\mathbf{t_2}[j]|}{\alpha_{2j}}\right)^{\beta_{2j}} - \left(\frac{|\mathbf{t_1}[j]|}{\alpha_{1j}}\right)^{\beta_{1j}}\right) \\
& + u(\mathbf{x})\left(\log\frac{1}{\sqrt{2\pi\sigma_{t1}^2}} - \log\frac{1}{\sqrt{2\pi\sigma_{t2}^2}}\right) + c,
\end{aligned} \tag{5.2.16}$$

with c independent of $u(.)$ or $\mathbf{t_i}$. This approach requires the actual generation of $g_1(.)$ and $g_2(.)$, which can be computationally demanding. Since many of the sparse texture representations correspond to a Parseval frame, this can also be evaluated in the transform domain:

$$\begin{aligned}
& \log \mathrm{p}_{\mathrm{F|U}}\big(f(\mathbf{x}) \mid u(\mathbf{x})\big) \\
= \;& u(\mathbf{x})\left(\frac{1}{2\sigma_{t2}^2}((S^{-1}f)(\mathbf{x}) - \mu_{2S})^2 - \frac{1}{2\sigma_{t1}^2}((S^{-1}f)(\mathbf{x}) - \mu_{1S})^2\right) \\
& + u(\mathbf{x})\left(\log\frac{1}{\sqrt{2\pi\sigma_{t1}^2}} - \log\frac{1}{\sqrt{2\pi\sigma_{t2}^2}}\right) + c,
\end{aligned} \tag{5.2.17}$$

with S^{-1} the inverse transform of S, i.e. $S\,S^{-1}\,f(.) = f(.)$. This analysis approach can be generalised to other transforms, e.g. to non-linear or non-invertible

transforms. This generalisation allows the use of many popular texture analysis transforms such as local binary patterns, Gabor filters, patch statistics (e.g. entropy and variance), principal curvature, intrinsic texture, morphological profiles, similarity with training texture, etc. [Rubner 00, Paragios 02, Yu 05, Houhou 08, Houhou 09, Babin 13b, Babin 13a, Orjuela-Vargas 14b, Orjuela-Vargas 14a]. Note that many of these transforms are not monomorphisms, i.e. two different texture patches can have the same transform output. This however does not have to be a drawback, as long as the texture of the foreground and the texture of the background results in a different output. Fig. 5.10 shows an example of texture transforms that are not monomorphisms. Fig. 5.10.B shows the entropy measured in a window around each pixel. Note that many texture patches can result in the same entropy, and that this transform cannot be inverted. However that does not hamper its usability for the detection of cell nuclei, which shows different values than the background. Fig. 5.10.C shows the result of a different popular texture transform: filtering with a Gabor filter. The generalisation of texture transforms to non monomorphisms results in the following log likelihood:

$$\begin{aligned} & \log \mathrm{p}_{\mathrm{F|U}}\big(f(\mathbf{x}) \mid u(\mathbf{x})\big) \\ = \; & u(\mathbf{x})\left(\frac{1}{2\sigma_2^2}((T\,f)(\mathbf{x})-\mu_{2T})^2 - \frac{1}{2\sigma_1^2}((T\,f)(\mathbf{x})-\mu_{1T})^2\right) \\ & + u(\mathbf{x})\left(\log\frac{1}{\sqrt{2\pi\sigma_{i1}^2}} - \log\frac{1}{\sqrt{2\pi\sigma_{i2}^2}}\right) + c, \end{aligned} \tag{5.2.18}$$

with T any specific texture transform. Note that some texture transforms are sensitive to scaling, small translations, rotation, etc., whereas other transforms are robust against (some of) these transformations. So extra care has to be taken in a suitable choice for T. For a detailed discussion on texture transforms and their invariance properties, we refer to [Orjuela-Vargas 13].

Dictionary based image models Similar to the texture models, the use of a linear combination of trained basis functions, i.e. a dictionary, is widely used for different image processing tasks, e.g. denoising, face recognition, super resolution, etc. [Yang 10b, Yang 10a, Tosic 11, Zeyde 12]. Inspired by the successes in these different applications, [Xia 12] proposed to define a foreground and background function, $g_1(.)$ and $g_2(.)$, as a linear combination of a set of dictionary functions, i.e. $g_1(\mathbf{x}) = \sum_j d_j(\mathbf{x})\,\mathbf{t_1}[j]$. These dictionary functions, $d_j(\mathbf{x})$, are calculated from a training dataset and are typically application dependent. These functions can either correspond to local intensity features such as edges or to local recurring intensity patches from the training dataset. For a general overview on the

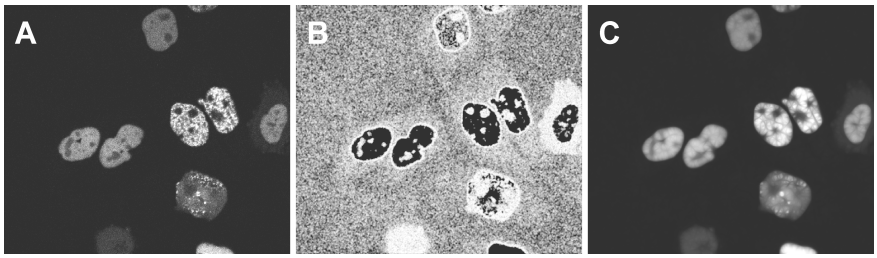


Figure 5.10: An example of texture in a fluorescence microscopic image. A) the original microscopic image, B) the entropy measured in a window around each pixel, C) the microscopic image filtered with a Gabor filter.

use and training of dictionaries in image processing we refer to [Rubinstein 10]. To prevent over-fitting, [Xia 12] proposed to model the likelihood of a coefficient of a specific dictionary function using a Laplacian distribution with mode zero. This results in the following log-likelihood for the image model

$$\begin{aligned} & \mathrm{p}_{\mathrm{G}_1(.)}(g_1(.))\,\mathrm{p}_{\mathrm{G}_2(.)}(g_2(.))\,\log \mathrm{p}_{\mathrm{F|U,G_1,G_2}}\big(f(\mathbf{x}) \mid u(\mathbf{x}), g_1(\mathbf{x}), g_2(\mathbf{x})\big) \\ = \;& u(\mathbf{x})\left(\frac{1}{2\sigma_2^2}\Big(f(\mathbf{x}) - \sum_i d_i(\mathbf{x})\,\mathbf{t_2}[i]\Big)^2 - \frac{1}{2\sigma_1^2}\Big(f(\mathbf{x}) - \sum_i d_i(\mathbf{x})\,\mathbf{t_1}[i]\Big)^2\right) \\ & -\frac{|\mathbf{t_2}|}{b_{2D}} - \frac{|\mathbf{t_1}|}{b_{1D}} + u(\mathbf{x})\left(\log\frac{1}{\sqrt{2\pi\sigma_{i1}^2}} - \log\frac{1}{\sqrt{2\pi\sigma_{i2}^2}}\right) + c, \end{aligned} \tag{5.2.19}$$

with $d_i(.)$ the i^{th} basis function of the dictionary and b_{iD} the diversity of the Laplace distribution used to model $\mathrm{p}_{\mathrm{G_i}(.)}(g_i(.))$.

5.2.2 A probabilistic prior

A correct image model generally allows distinguishing segment pixels from background pixels. Under the influence of noise, clutter and blur, the image model sometimes result in a similar likelihood for both segment and background. In such a situation, the segmentation accuracy of the MAP estimator will significantly increase if a proper prior is used, i.e. prior knowledge about the object that is to be segmented. Literature contains an abundance of segmentation priors, ranging from general-purpose priors to application specific shape priors, e.g. modelling tree structures [Stuhmer 13], the area ratio between objects [Nieuwenhuis 13b], the proximity of certain objects [Bergbauer 13], etc. In this section we will elaborate on general-purpose priors, ranging from generic priors to full shape modelling priors.

5.2.2.1 Perimeter & area

While many shape priors exist for specific applications, simple but generic priors often already result in accurate segmentation. Common generic shape priors use the area, m_1, and perimeter, m_2, of the detected objects [Mumford 89, Chan 99]. For a continuous characteristic function, i.e. a binary function, these measurements can respectively be expressed as $m_1(u(.)) = |u(.)|$ and $m_2(u(.)) = |\nabla u(.)|$ (see [Evans 92] for more details). We propose to model the probability of the area-perimeter of a segment using a generalised Laplacian, i.e.

$$\mathrm{p}(m_i(u(.))) = \frac{\beta_{mi}}{2\alpha_{mi}\Gamma(\frac{1}{\beta_{mi}})} e^{-\left(\frac{|m_i(u(.))-\mu_{mi}|}{\alpha_{mi}}\right)^{\beta_i}}, \tag{5.2.20}$$

with μ_1 the expected area, μ_2 the expected perimeter, and α_{mi}, β_{mi} trained fitting parameters. For simplicity we assume that area and perimeter are independent of each other. This results in the following segmentation prior:

$$\begin{aligned} \log \mathrm{p}(u(.)) \quad = \quad & -\left(\frac{|\,|u(.)| - \mu_{m1}|}{\alpha_{m1}}\right)^{\beta_{m1}} + \log\left(\frac{\beta_{m1}}{2\alpha_1\Gamma(\frac{1}{\beta_{m1}})}\right) \\ & -\left(\frac{||\nabla u(.)| - \mu_{m2}|}{\alpha_{m2}}\right)^{\beta_{m2}} + \log\left(\frac{\beta_{m2}}{2\alpha_{m2}\Gamma(\frac{1}{\beta_{m2}})}\right). \end{aligned} \tag{5.2.21}$$

Note that since $u(\mathbf{x})$ is positive for all $\mathbf{x}$ that

$$\left(\frac{|\,|u(.)| - \mu_{m1}|}{\alpha_{m1}}\right)^{\beta_{m1}} = \left(\frac{|u(.) - \mu_{m1}|}{\alpha_{m1}}\right)^{\beta_{m1}} \tag{5.2.22}$$

, which results in a convex function. However the term related to the total variation is not necessarily convex, depending on the actual parameters of the generalised Laplacian. For β_{m1} and β_{m2} equal to 1, this prior corresponds to original Mumford-Shah constraints [Mumford 89], while β_{m1} equal to 1, β_{m2} equal to 0 in combination with the piecewise constant Mumford-Shah image model (section 5.2.1.1) results in the convex active contour model from section 5.1.2.

5.2.2.2 Markov random fields

Note that if for a certain pixel, all its neighbours are classified as segment pixels, that the pixel itself is likely a segment pixel as well. This concept can be further generalised: the likelihood that a pixel belongs to the segment, depends on the classification of the other pixels. This requires the modelling of correlation

between pixels. To simplify this, a Markov assumption is often used, i.e. the assumption that the likelihood of $u(\mathbf{x})$ only depends on the likelihood of pixels in the neighbourhood of that pixel, $\Omega_{\mathbf{x}}$.

The exact definition of a neighbourhood can depend on the application, but a common definition results in all pixels within a predefined distance, r, i.e.

$$\Omega_{\mathbf{x}} = \{\mathbf{y} \in \mathbb{Z}^2 \backslash \mathbf{x} \mid \|\mathbf{x} - \mathbf{y}\| \leq r\}. \tag{5.2.23}$$

Fig. 5.11.A shows an example of such a neighbourhood for $r = \sqrt{2}$. This definition of a neighbourhood allows us to formalise the Markov property:

$$\mathrm{p}(u(\mathbf{x}) \mid u(\mathbf{y})_{\mathbf{y} \in \mathbb{Z}^2 \backslash \mathbf{x}}) = \mathrm{p}(u(\mathbf{x}) \mid u(\mathbf{y})_{\mathbf{y} \in \Omega_{\mathbf{x}}}), \tag{5.2.24}$$

i.e. knowing the perfect segmentation output of all pixels except for the pixel $\mathbf{x}$, gives no more information from knowing the segmentation output for all the pixels in the neighbourhood, $\Omega_{\mathbf{x}}$, with exception of $\mathbf{x}$. Irrespective of which pixels belong to a neighbourhood, it is also useful to define which pixels are direct neighbours, i.e. which pixels are connected. The most common connectivities for two dimensional images are the 4- and 8-connectivities. The direct neighbours for a pixel $\mathbf{x}$, i.e. the pixels that are connected to $\mathbf{x}$, are defined as:

$$\Theta_{\mathbf{x}} = \{\mathbf{y} \in \mathbb{Z}^2 \mid \|\mathbf{x} - \mathbf{y}\|_1 = 1\} \tag{5.2.25}$$

for the 4-connectivity, and

$$\Theta_{\mathbf{x}} = \{\mathbf{y} \in \mathbb{Z}^2 \mid \|\mathbf{x} - \mathbf{y}\|_\infty = 1\} \tag{5.2.26}$$

for the 8-connectivity, with $\|.\|_1$ the l_1 or Manhattan norm and $\|.\|_\infty$ the l_∞ or supremum norm. Fig. 5.11.B shows the direct neighbours of the central pixel of Fig. 5.11.A for both 4- and 8-connectivity.

A clique C of a neighbourhood, $\Omega_{\mathbf{x}}$, is a subset of pixels in $\Omega_{\mathbf{x}}$, such that $\mathcal{C}$ is either a singleton, a pair direct neighbours or a collection of n direct neighbouring pixels. The number of pixels within a clique, is called the order of the clique. The set of all cliques, $\mathcal{C}$ is then defined as follows:

$$\begin{aligned} \mathcal{C} &= \{\mathbf{y} \in \Omega_{\mathbf{x}}\} \\ &\cup \{(\mathbf{y}_1, \mathbf{y}_2) \in \Omega_{\mathbf{x}} \times \Omega_{\mathbf{x}} \mid \mathbf{y}_1 \in \Theta_{\mathbf{y}_2}\} \\ &\cup \{(\mathbf{y}_1, \mathbf{y}_2, \mathbf{y}_3) \in \Omega_{\mathbf{x}}^3 \mid \mathbf{y}_1 \in \Theta_{\mathbf{y}_2}, \mathbf{y}_2 \in \Theta_{\mathbf{y}_3}\} \\ &\cup \ldots \end{aligned} \tag{5.2.27}$$

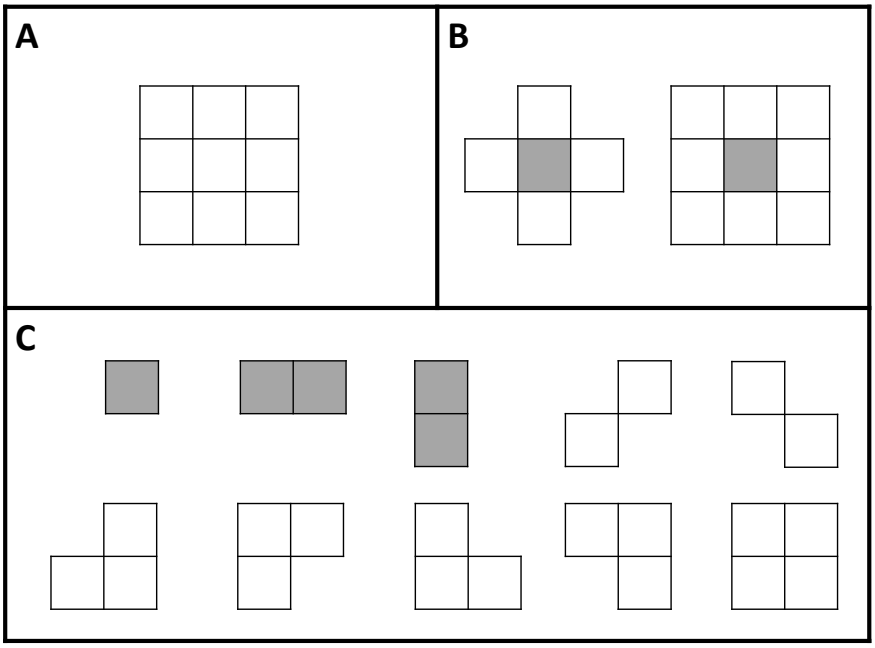


Figure 5.11: A) An example of a neighbourhood, B) the direct neighbours of the central pixel of the neighbourhood assuming 4-connectivity on the left, and 8-connectivity on the right. C) The type of cliques up to order two for the neighbourhood of A. The grey cliques correspond to 4-connectivity, whereas 8-connectivity contains both grey and white type of cliques.

According to the Hammersley-Clifford theorem [Clifford 90], the joint probability of $u(.)$, which is a global prior, can be expressed using only local clique potentials:

$$\mathrm{p}(u(.)) = \frac{1}{\alpha} e^{-\frac{1}{\beta}\sum_{c\in\mathcal{C}} V_c(u(.))}, \tag{5.2.28}$$

where α a normalising coefficient, β a parameter manipulating the kurtosis of this Gibbs distribution, and $V_c(.)$ represents the clique potential representing the interaction of neighbouring labels over the clique. A simple example of a Markov shape prior is the Potts model [Potts 52, Nieuwenhuis 13c], where only pairwise cliques are considered and the clique potentials are defined as:

$$V(u(\mathbf{x}),\, u(\mathbf{y})) = \begin{cases} \gamma & \text{if } u(\mathbf{x}) \neq u(\mathbf{y}) \\ 0 & \text{if } u(\mathbf{x}) = u(\mathbf{y}) \end{cases}. \tag{5.2.29}$$

The use of other cliques and clique potentials has led to accurate segmentation of anatomical structures [Despotovic 14], individual cell nuclei in microscopic images [Nemeth 11], tree crowns in aerial images [Horvath 09], etc. For a more in depth overview of the use of MRF in image processing we refer to [Li 09, Blake 11, Ružić 14].

5.2.2.3 Shape descriptors

With the introduction of active contours with convex energies (sec. 5.1.2), the proposed segmentation methods became invariant to the initialisation of $u(.)$. However, in situations where a good initialisation is available, e.g. from a previous frame in a video sequence, it is desirable to have a final segmentation similar to the initialisation, even if this segmentation corresponds to a slightly lower likelihood. The parametric active contours achieve this by optimising to local optima in a non-convex energy function (see chapter 4 for more details). It is also possible to achieve such a similarity with convex terms. The first step is to transform $u(.)$ to a small set of shape descriptors [Cremers 08, Klodt 13]. Especially linear transforms are useful since they result in convex terms. Plenty of shape descriptors based on linear transforms have been proposed in literature: eigenfaces, geometric moments, Zernike moments, Fourier-Mellin descriptors, etc. For a detailed overview of shape descriptors we refer to chapter 8. For simplicity we assume that most of the spatial correlation is removed by the shape transform, and as such that all shape descriptors are uncorrelated. This allows the use of a generalised Laplacian to model the shape prior, i.e.

$$\mathrm{p}(u(.)) = \frac{\beta_{\boldsymbol{s}}}{2\alpha_s\Gamma(\frac{1}{\beta_s})} e^{-\left(\frac{|Su(.)-\boldsymbol{\mu_s}|}{\alpha_i}\right)^{\beta_i}}, \tag{5.2.30}$$

with S a shape transform of choice, $\beta_s > 1$ the shape parameter of the generalised Laplacian distribution, α_s the scale parameter and $\boldsymbol{\mu_s}$ the expected values for the shape descriptors, e.g. estimated from the segmentation result from previous frames. As an example, consider the work of Klodt et al. [Klodt 11, Klodt 13], where the geometric moments were used as a shape transform. In this case, modelling only the geometric moments up to the second order, result in penalising segmentation results that have not the same centroid, scale and orientation as a certain input shape.

5.3 Optimisation

Several methods have been proposed to optimise geometric active contours [Boyd 04, Chan 06, Bresson 07b, Yip 09, Goldstein 10, Chen 12]. Although these methods all claim fast convergence, they have not been properly compared. In this section we will first summarise some of these methods before experimentally evaluating the convergence and constraints of these methods. As a use case we will optimise the popular constant Mumford-Shah model with perimeter regularisation, i.e. we will optimise the convex energy defined in eq. (5.1.7).

5.3.1 Gradient descent

One of the simplest optimisation schemes minimises eq. (5.1.7) by solving the following set of Euler-Lagrange equations using gradient descent:

$$\nabla \cdot \Big(\frac{\nabla u(\mathbf{w})}{|\nabla u(\mathbf{w})|} \Big) - \beta\, r(\mathbf{w}) = 0, \tag{5.3.1}$$

with $\mathbf{w}$ a place vector, and $r(\mathbf{w}) = (f(\mathbf{w}) - \mu_1)^2 - (f(\mathbf{w}) - \mu_2)^2$. In order to constrain the solution of this equation to the interval $[0, 1]$ a convex barrier or potential function, $b(.)$, is added as an extra term to eq. (5.1.7), i.e.

$$E[u(.)] = |\nabla u(.)| + \beta \langle u(.), r(.) \rangle + b(u(.)). \tag{5.3.2}$$

A wide variety of barrier functions have been proposed in literature, e.g. quadratic, linear, Huber, Lipschitz functions etc. Fig. 5.12 shows four different barrier functions for the interval $[0, 1]$. For a more extensive review on different barrier functions, we refer to [Pock 10]. For this specific application, we use the piecewise linear barrier function, since it was proven in [Chan 06] that this function enforces $u(\mathbf{w}) \in [0, 1]$ exactly. And as can be seen in Fig. 5.12, this function does not penalize or favour any value within $[0, 1]$. This piecewise linear barrier function can be defined as

$$b(x) = \max \Big(0, 2\,\gamma \Big| x - \frac{1}{2} \Big| - 1 \Big), \tag{5.3.3}$$

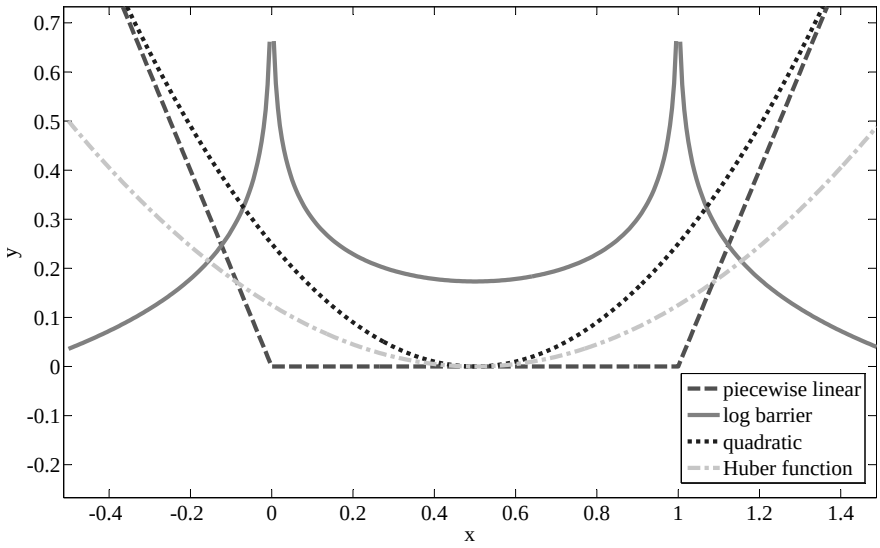


Figure 5.12: Four different examples of a barrier function that can be used to keep optimisers within the interval $[0, 1]$ [De Vylder 11a].

where γ is the maximum intensity of the image. Adding this penalty function to eq. (5.3.1) results in the following gradient descent optimisation step:

$$u^{(k+1)}(\mathbf{w}) = u^{(k)}(\mathbf{w}) + \delta t\, \phi_{gd}(u^{(k)}(\mathbf{w})) - \delta t \frac{d\, b(u^{(k)}(\mathbf{w}))}{d\, u^{(k)}(\mathbf{w})}, \tag{5.3.4}$$

where δt is the stepsize, and

$$\phi_{gd}(u^{(k)}(\mathbf{w})) = \frac{\partial\Big(\frac{\frac{\partial u^{(k)}(\mathbf{w})}{\partial x}}{\|\nabla u^{(k)}(\mathbf{w})+\epsilon\|}\Big)}{\partial x} + \frac{\partial\Big(\frac{\frac{\partial u^{(k)}(\mathbf{w})}{\partial y}}{\|\nabla u^{(k)}(\mathbf{w})+\epsilon\|}\Big)}{\partial y} - \beta\, r(\mathbf{w}). \tag{5.3.5}$$

Here ϵ represents a small constant, added in order to compensate for the non-differentiability of the total variation norm.

5.3.2 Newton-Raphson

A similar approach to solve the Euler-Lagrange equations in eq. (5.3.1) is done by iteratively updating according to the Newton-Raphson optimisation scheme, i.e.

$$u^{(k+1)}(\mathbf{w}) = u^{(k)}(\mathbf{w}) + \frac{\phi_{gd}(u^{(k)}(\mathbf{w}))}{\phi_{nr}(u^{(k)}(\mathbf{w}))}, \tag{5.3.6}$$

where $\phi_{nr}(.)$ is the derivative of $\phi_{gd}(.)$, i.e.

$$\begin{aligned}\phi_{nr}(u^{(k)}(\mathbf{w})) &= \frac{d\,\phi_{gd}(u^{(k)}(\mathbf{w}))}{d\,u^{(k)}(\mathbf{w})} \\ &= 2\Big(\frac{-2}{\|\nabla u^{(k)}(\mathbf{w})+\epsilon\|} + \frac{(\nabla\cdot u^{(k)}(\mathbf{w}))^2}{\|\nabla u^{(k)}(\mathbf{w})+\epsilon\|^3}\Big).\end{aligned}$$

The Newton-Raphson optimisation scheme should theoretically converge faster than gradient descent, i.e. it has a quadratic convergence rate. In sec. 5.3.7 we experimentally validate the actual convergence rate.

5.3.3 Line search

In the gradient descent method the variable $u(\mathbf{w})$ is updated according to the gradient direction with a constant step size δt. A variation on this optimisation scheme is given by a line search algorithm, where the updating direction is the same as with gradient descent, but where the step size is variable. We propose a line search like optimisation approach, where the step size is chosen such that the value of $u(\mathbf{w}) \in \{0,1\}$. This is a reasonable assumption, since a thresholded version of an optimal $\hat{u}(\mathbf{w})$ is also optimal (see section 5.1.2 or [Bresson 07a]). So instead of looking for an optimal $\hat{u}(\mathbf{w})$ and then threshold it, we will immediately search for a binary optimum. This is achieved by the following updating scheme:

$$u^{(k+1)}(\mathbf{w}) = u^{(k)}(\mathbf{w}) + \phi_{ls}(u^{(k)}(\mathbf{w}))\,\mathrm{sgn}(\phi_{gd}(u^{(k)}(\mathbf{w}))), \tag{5.3.7}$$

where $\phi_{ls}(.)$ represents the step size, i.e.

$$\phi_{ls}(u^{(k)}(\mathbf{w})) = \begin{cases} 0 & \text{if } u^{(k)}(\mathbf{w}) = 0 \text{ and } \phi_{gd}(u^{(k)}(\mathbf{w})) < 0 \\ 0 & \text{if } u^{(k)}(\mathbf{w}) = 1 \text{ and } \phi_{gd}(u^{(k)}(\mathbf{w})) > 0 \\ 1 & \text{otherwise.} \end{cases} \tag{5.3.8}$$

This step size results in an optimisation step where $u^{(k+1)}(\mathbf{w})$ either keeps the same value or where it switches its value. Assume for example that $u^{(k)}(\mathbf{w}) = 0$ and that the gradient of eq. (5.3.1) points in the direction of $-\infty$, then $u^{(k+1)}(\mathbf{w})$ will remain zero. If the gradient pointed in the direction of $+\infty$ however, $u^{(k+1)}(\mathbf{w})$ would switch to the value one. Note that this method is more a heuristic than a real optimisation technique, i.e. there is no proof of convergence. In section 5.3.7 the convergence results of this method will be tested and discussed.

5.3.4 Dual formulations

A different approach is proposed in [Bresson 07b, Yip 09]. The gradient-based approaches typically require a small timestep in order to avoid numerical instabil-

ity caused by the total variation norm. To circumvent this slow convergence, an optimisation scheme is calculated based on a dual formulation of the total variation norm [Aujol 05]. Instead of optimising eq. (5.3.2) directly for $u(.)$, an extra function $v(.)$ is introduced in the energy term, i.e. the following energy term is optimised for $u(.)$ and $v(.)$:

$$E[u(.), v(.)] = |\nabla u(.)| + \beta\langle v(.),\, r(.)\rangle + \frac{\lambda}{2}\|u(.) - v(.)\|^2 + b(v(.)), \quad (5.3.9)$$

with $b(.)$ the same barrier function as in subsection 5.3.1 and λ a weighting parameter, defining the influence of the similarity between $u(.)$ and $v(.)$. Due to the convexity of the energy function, we can alternate between the following two optimisation problems:

$$\arg\min_{u(.)} \quad |\nabla u(.)| + \frac{\lambda}{2}\|u(.) - v(.)\|^2 \quad (5.3.10)$$

$$\arg\min_{v(.)} \quad \beta\langle v(.),\, r(.)\rangle + \frac{\lambda}{2}\|u(.) - v(.)\|^2 + b(v(.)). \quad (5.3.11)$$

There are numerous generic optimisation methods for solving $l_1 - l_2$ equations such as eq. (5.3.10) [Zibulevsky 10, Chambolle 11]. However, for this specific equation more efficient methods can be used. It can be proven that $u^{(k+1)}(.)$ satisfies the following optimality condition [Bresson 07b] :

$$u^{(k+1)}(\mathbf{w}) = v^{(k)}(\mathbf{w}) - \frac{1}{\lambda}\nabla \cdot p(\mathbf{w}), \quad (5.3.12)$$

where $p(.)$ satisfies

$$\nabla\big(\frac{1}{\lambda}\nabla \cdot p(\mathbf{w}) - v^{(k)}(\mathbf{w})\big) - |\nabla\big(\frac{1}{\lambda}\nabla \cdot p(\mathbf{w}) - v^{(k)}(\mathbf{w})\big)|\, p(\mathbf{w}) = 0. \quad (5.3.13)$$

Given this set of equations, $p(.)$ can be efficiently calculated using the following fixed point algorithm [Aujol 05]:

$$\begin{aligned} p^{(0)}(\mathbf{w}) &= (0,0) \\ p^{(l+1)}(\mathbf{w}) &= \frac{p^{(l)}(\mathbf{w}) + \delta t\, \nabla(\nabla \cdot p^{(l)}(\mathbf{w}) - \lambda v^{(k)}(\mathbf{w}))}{1 + \delta t|\nabla(\nabla \cdot p^{(l)}(\mathbf{w}) - \lambda v^{(k)}(\mathbf{w}))|}. \end{aligned}$$

Experiments show that a single iteration of this fixed point method already results in good convergence properties for the optimisation of eq. (5.3.10) and by extension for the optimisation of eq. (5.3.9). The optimisation of $v(.)$ in eq. (5.3.11) is give by the following closed solution:

$$v^{(k+1)}(\mathbf{w}) = \min\Big(\max\big(u^{(k+1)}(\mathbf{w}) - \frac{\beta}{\lambda}r(\mathbf{w}), 0\big), 1\Big). \quad (5.3.14)$$

For the proof of this closed solution, we refer to [Bresson 07b].

5.3.5 Split alternating direction method of multipliers

A different optimisation method is the split alternating direction method of multipliers optimisation (ADMM), which is similar to the augmented Lagrangian method and is considered to be an efficient optimisation technique for solving l_1 regularised problems [Goldstein 09, Goldstein 10, Mao 10]. Similar to the optimisation based on the dual formulation, the split ADMM method will "de-couple" the l_1 norm and the inner product, by introducing a new function $d(.)$ and by putting constraints on this new variable. This results in the following optimisation problem:

$$E[u(.),\, d(.)] = |d(.)| + \beta\langle u(.),\, r(.)\rangle \text{ such that } d(.) = \nabla u(.). \tag{5.3.15}$$

This optimisation problem can be converted to an unconstrained problem by adding a quadratic penalty function, i.e.

$$E[u(.),\, d(.)] = |d(.)| + \beta\langle u(.),\, r(.)\rangle + \frac{\lambda}{2}\|d(.) - \nabla u(.)\|^2, \tag{5.3.16}$$

where λ is a weighting parameter. Setting λ high introduces numerical instability, whereas when λ is not high, the quadratic penalty function only approximates the constraint $d(.) = \nabla u(.)$. This is typically overcome by using different λ for each iteration. ADMM provides a different way of optimising 5.3.15, while enforcing the constraint exactly in an efficient way [Goldstein 09]. In the ADMM technique an extra function, $b^{(k)}(.)$ is subtracted from the penalty function. This results in the following two unconstrained steps:

$$\begin{aligned}(u^{(k+1)}, d^{(k+1)}) &= \underset{u^{(k)}(.), d^{(k)}(.)}{\arg\min}\ |d^{(k)}(.)| + \beta\langle u^{(k)}(.),\, r(.)\rangle \\ &\qquad + \frac{\lambda}{2}\|d^{(k)}(.) - \nabla u^{(k)}(.) - b^{(k)}(.)\|^2 \qquad (5.3.17)\\ b^{(k+1)}(.) &= b^{(k)}(.) + \nabla u^{(k+1)}(.) - d^{(k+1)}(.). \qquad (5.3.18)\end{aligned}$$

The first step requires optimising for two different functions, $u(.)$ and $d(.)$. Since the constrained optimisation step is convex, these optimal vectors can be calculated by alternating between optimising eq. (5.3.18) for $u(.)$ and optimising for $d(.)$:

$$\begin{aligned}u^{(k+1)}(.) &= \underset{u^{(k)}(.)}{\arg\min}\, \beta\langle u^{(k)}(.),\, r(.)\rangle + \frac{\lambda}{2}\|d^{(k)}(.) - \nabla u^{(k)} - b^{(k)}\|^2 \\ d^{(k+1)}(.) &= \underset{d^{(k)}(.)}{\arg\min}|d^{(k)}(.)| + \frac{\lambda}{2}\|d^{(k)}(.) - \nabla u^{(k+1)} - b^{(k)}\|^2. \qquad (5.3.19)\end{aligned}$$

The first problem can be optimised by solving a set of Euler-Lagrange equations. For each element $u(\mathbf{w})$ the following optimality condition should be satisfied:

$$\nabla^2 u^{(k)}(\mathbf{w}) = \frac{\beta}{\lambda}\, r(\mathbf{w}) + \nabla \cdot (d^{(k)}(\mathbf{w}) - b^{(k)}(\mathbf{w})). \tag{5.3.20}$$

In [Goldstein 10] it was proposed to solve this linear system using the iterative Gauss-Seidel method. However, in order to guarantee convergence of the Gauss-Seidel method, the system matrix should be strictly diagonally dominant or should be positive definite. Unfortunately neither is the case for eq. (5.3.20). Therefore we will optimize eq. (5.3.20) using the iterative conjugate residual method, which is a Krylov subspace method for which convergence is guaranteed if the system matrix is Hermitian [Saad 03]. Experiments showed that due to the fact that this system is not constrained to $[0, 1]$, too many iterations of the conjugate residual method results in numerical unstable solutions. In [Goldstein 10] it was observed that the ADMM optimisation converges even with just an approximation of the solution of the linear system in eq. (5.3.20). Therefore, we propose to use only one iteration of the conjugate residual method. Using just one iteration gives a reduced algorithm that needs less calculations and less memory:

$$\begin{aligned} c(\mathbf{w}) &= \frac{\beta}{\lambda} r(\mathbf{w}) + \nabla \cdot (d^{(k)}(\mathbf{w}) - b^{(k)}(\mathbf{w})) - \nabla^2 u^{(k)}(\mathbf{w}) \\ \alpha &= \frac{\langle c(.),\, \nabla^2 c(.) \rangle}{\langle \nabla^2 c(.),\, \nabla^2 c(.) \rangle} \\ u_u^{(k+1)}(\mathbf{w}) &= u^{(k)}(.) - \alpha c(\mathbf{w}). \end{aligned}$$

The solution of this conjugate residual is unconstrained, i.e. $u_u^{(k+1)}(\mathbf{w})$ does not have to lie in the interval $[0, 1]$. Note that minimising eq. (5.3.19) for $u_u^{(k+1)}(\mathbf{w})$, is equivalent to minimising a quadratic function. If $u_u^{(k+1)}(\mathbf{w}) \notin [0, 1]$ then the constrained optimum is either 0 or 1, since a quadratic function is monotonic in an interval which does not contain its extremum. So the constrained optimum is given by:

$$u^{(k+1)}(\mathbf{w}) = \max\Big(\min\big(u_u^{(k+1)}(\mathbf{w}), 1\big), 0\Big). \tag{5.3.21}$$

In order to calculate an optimal $d^{(k+1)}(.)$ in eq. (5.3.19), a closed form solution can be calculated using the shrinking operator, i.e.

$$d^{(k+1)}(\mathbf{w}) = \text{shrink}\Big(\nabla u_u^{(k+1)}(\mathbf{w}) + b^{(k)}(\mathbf{w}), \lambda\Big), \tag{5.3.22}$$

where

$$\text{shrink}(\tau, \lambda) = \begin{cases} 0 & \text{if } |\tau| \leq \frac{1}{\lambda} \\ \tau - \frac{1}{\lambda}\,\text{sgn}(\tau) & \text{otherwise.} \end{cases} \tag{5.3.23}$$

5.3.6 Other methods

This enumeration of optimisation methods is obviously not exhaustive. Although many other optimisation methods have been mentioned in literature, we will not

discuss them in detail. We will however briefly give an overview of other groups of optimisation approaches:

Graph based approaches: By representing the images as a graph where each pixel corresponds to a node and the vertices connect neighbouring pixels, segmentation is equivalent to partitioning the graph in order to minimise an energy function. Literature reports on efficient partitioning algorithms, such as graph-cuts, that result in good segmentation [Lempitsky 07, Vicente 08, Despotovic 14]. For a comparison between graph-cuts optimisation and some of the optimisation methods discussed in this chapter, we refer to [Bresson 09].

Primal-dual approaches: In these type of approaches, the original optimisation problem (the primal problem) can be transformed to another optimisation problem (the dual problem), where the optimal solution is equal for both primal and dual problem. Typical dual problems are calculated based on the Lagrangian or on the Legendre-Fengel transform. Efficient optimisation is achieved by iteratively minimising the difference between the solution of the primal problem and the solution of the dual problem, i.e. the so called duality gap [Yip 09, Krishnan 07, Chambolle 11, Zhang 11].

Meta-heuristics: Techniques such as simulated annealing and swarm optimisation have successfully been used for the optimisation of geometric active contours [Horvath 06, Horvath 09, Nemeth 11]. These are especially interesting for non-convex energies, but their lack of guarantee of finding the optimum makes them less desirable for active contours with convex energy terms.

Hybrid methods: Although heuristical methods, such as simulated annealing or the proposed line search method, lack theoretical convergence properties, they do have their merit. In order to guarantee convergence to a global optimum, optimisation methods can be combined: first a heuristical approach can be used to get a good approximation of the optimum, this approximation can then be used as an initialisation for other optimisation methods.

For a more in depth discussion on optimisation we refer to [Boyd 04].

5.3.7 Experimental results

For the validation of the segmentation, a small dataset of 7 pictures was assembled. These pictures come from different image modalities such as microscopy, photos, ultrasound and MRI-scans. All these images have a dimension of 512×512 pixels.

To compare the active contour result with ground truth segmentation, the Dice coefficient is used (see Appendix B). Note that the Dice coefficient is a relative measurement, i.e. the segmentation result is compared to ground truth. Ground truth however is a relative concept. What is considered to be a foreground segment for one application might be clutter for a different application. In this section we will compare different optimisation methods, and thus focus on how the methods converge to a certain segmentation result instead of focussing on the quality of the actual MAP estimator. As such, we will compare the calculated active contour with the optimal contour, i.e. the contour with minimal energy. Therefore, the segmentation result is calculated using all optimisation methods described in this section. The active contour which had the lowest energy after 300 iterations is used as ground truth.

5.3.7.1 Convergence

As a first experiment, each method runs 300 iterations on the test dataset. After each iteration, the Dice coefficient between the thresholded active contour and the ground truth was calculated. The left part of Fig. 5.13 shows the average Dice coefficient after each iteration for each method. As can be seen, the gradient descent method is the slowest, this could be improved by changing the stepsize, although this risks missing the actual optimum. All other methods converge fast during the first few iterations, but then slow down. The right part of Fig. 5.13 shows a close up of the left plot. The line search converges much faster during the first iteration than other methods, but then starts oscillating, and never really converges. Optimisation based on the dual formulation of the TV-norm is the second fastest method and does converge to a global optimum. Both the Newton-Raphson and the ADMM method converge slower than previous mentioned methods. Neither reach the real optimum due to numerical errors. In order to benefit from the fast start from the line search like optimisation, we propose a hybrid optimisation with a single updating step based on line search and for all other iterations applying Newton-Raphson iterations. This combined optimisation reaches the real optimum, while converging faster than all other methods.

Since a higher convergence rate not necessarily means faster optimisation, we summarize some extra relevant measurements in Table 5.1. The second column of this table contains the theoretical time complexity of a single iteration for each method. In this setting we assume that an image has a dimension of $n \times m$. The dual formulation optimisation requires the calculation of the following images of size $n \times m$: the calculation of $p(.)$, the calculation of $u(.)$ and finally the calculation of $v(.)$. The ADMM method requires iterating 9 times over each pixel of an $n \times m$ image: 7 runs in the conjugate residual optimisation and two runs to update the functions $d(.)$ and $b(.)$. The conjugate residual needs a single run to calculate the right side

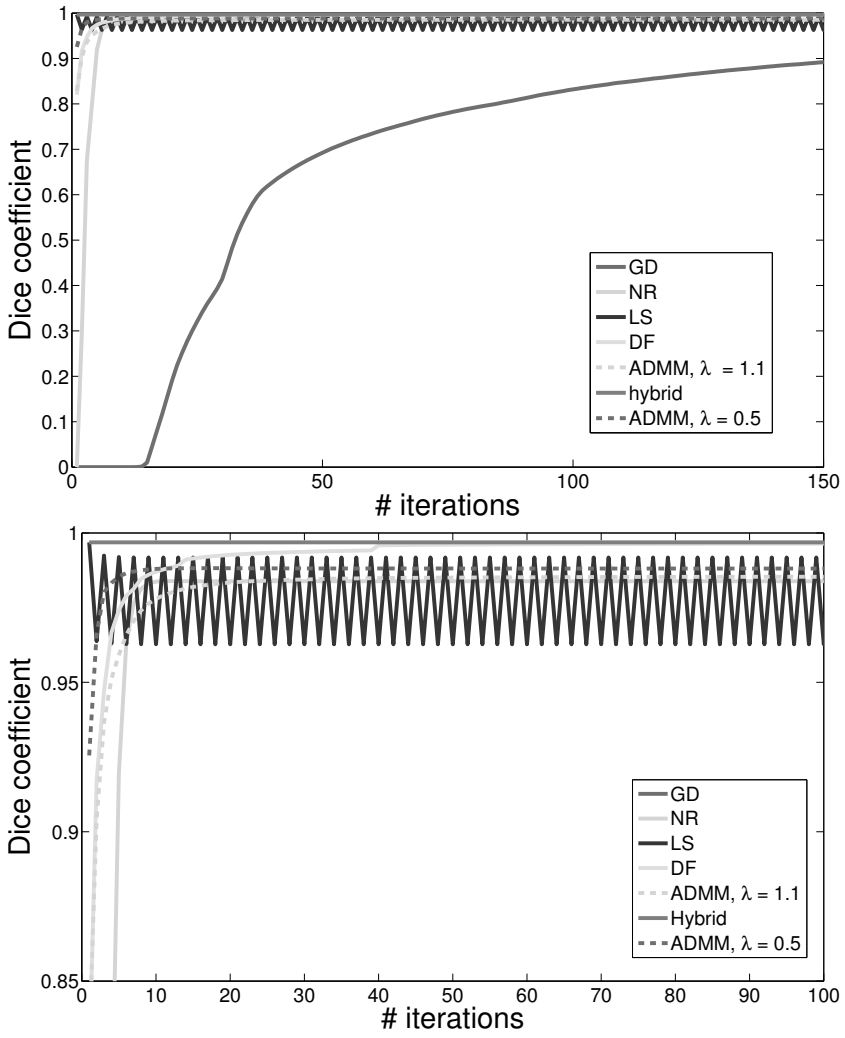


Figure 5.13: A comparison of the convergence of the Dice coefficient between ground truth and the active contours, calculated using different methods. The methods used for the active contour optimisation are: gradient descent, Newton-Raphson, line search, dual formulation, ADMM and a combination of line search and Newton-Raphson. Below a close up of the top graph is shown [De Vylder 11a].

Table 5.1: A comparison of different optimisation methods.

method	time (complexity)	time (ms)	memory
gradient descent	$O(n \times m)$	30.83	$O(1)$
Newton-Raphson	$O(n \times m)$	33.91	$O(1)$
Line search	$O(n \times m)$	25.79	$O(1)$
hybrid	$O(n \times m)$	32.49	$O(1)$
Dual formulation	$O(3 \times n \times m)$	33.80	$O(3 \times n \times m)$
ADMM	$O(9 \times n \times m)$	60.09	$O(7 \times n \times m)$

of the Euler-Lagrange equations, two runs to calculate a Laplacian, one run to calculate the residuals, two runs to calculate an inner product and finally a single run to update $u(.)$.

The third column of Table 5.1 contains the average time to calculate a single iteration of a 512×512 image.[1] The ADMM has the highest computational cost, as in correspondence with the theoretical complexity of the method. All other methods result in a similar computational time, although the line search method is slightly fastest to calculate. The last column summarizes the memory required for each method. We have omitted the memory needed to store the input image and the actual segmentation $u^{(k)}(.)$, since this is the same for each method. Only the dual formulation method and the ADMM method require more memory to store extra variables. The dual formulation method needs to store $v(.)$ and $p(.)$, $v(.)$ has the same size as the image, whereas $p(.)$ is twice as big as the image. For the ADMM the extra functions $d(.)$ and $b(.)$ have to be stored, each twice the size of the image. The conjugate residual method, needed for the optimisation of the ADMM method, requires three times the amount of memory to store the image. This is needed to store the residuals, the Laplacian of the contour itself and the Laplacian of the residuals. The Laplacians could be recalculated, instead of stored, but such an implementation would lead to slower iterations.

5.3.7.2 Scalability

To test the scalability of the different optimisation algorithms, the average calculation time of a single iteration was measured for images of different sizes. These images have a dimension of $N \times N$ where $N = 2^7, 2^8, ..., 2^{12}$. Fig. 5.14 shows the results of these time measurements. The ADMM is generally slower than the other methods. The line search method remains the fastest to calculate, independent of the image size. The combination of the line search method and the Newton-Raphson method is slower than optimisation only based on the line search

[1] All methods were implemented in C and run on an Intel i7 Q720 1.6 GHz CPU with 4GB RAM.

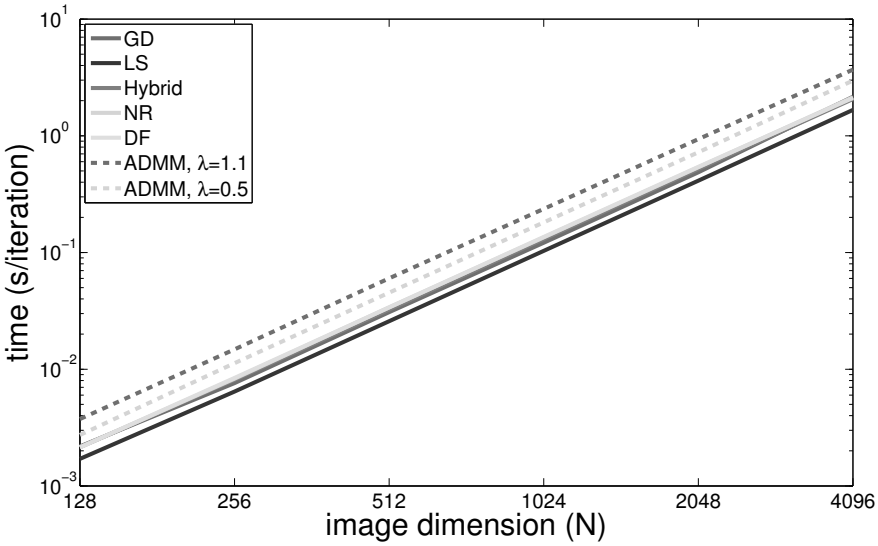


Figure 5.14: The time needed to calculate a single optimisation iteration, depending on the image size. The images have dimension $N \times N$.

scheme, but is still faster to calculate than the other discussed methods. Aside from computational speed, the memory requirements are an extra factor determining the scalability. Although the dual formulation method is reasonably fast for large 2D images, e.g. images of 4096×4096 pixels, it does require three times the amount of memory needed by the gradient-based methods. This memory constraint will hamper the method for even bigger datasets, e.g. for segmentation of 3D images or for high resolution images such as from ultra structures images with scanning electron microscopy (SEM), e.g. images up to $32K \times 32K$.

5.3.7.3 Robustness to noise

The experiments in subsection 5.3.7.2 are all done using the same parameters. Although the measurements are done on different types of images, e.g. microscopy, photos, etc., these measurements say little on the influence of noise or on the influence of the regularisation term. To test these influences, we added white Gaussian noise to the images ($\sigma^2 = 2.56$, for images with an intensity between 0 and 355). The active contours were optimised with a small weight for the data fit term, i.e. $\beta = 0.0001$ in eq. (5.1.9), thus emphasising the influence of the total variation regularisation term. A new set of ground truth images was calculated for these noisy images. Fig. 5.15 shows the results of the convergence measured with the Dice coefficient. Just as with the non-noisy images, both the dual formulation

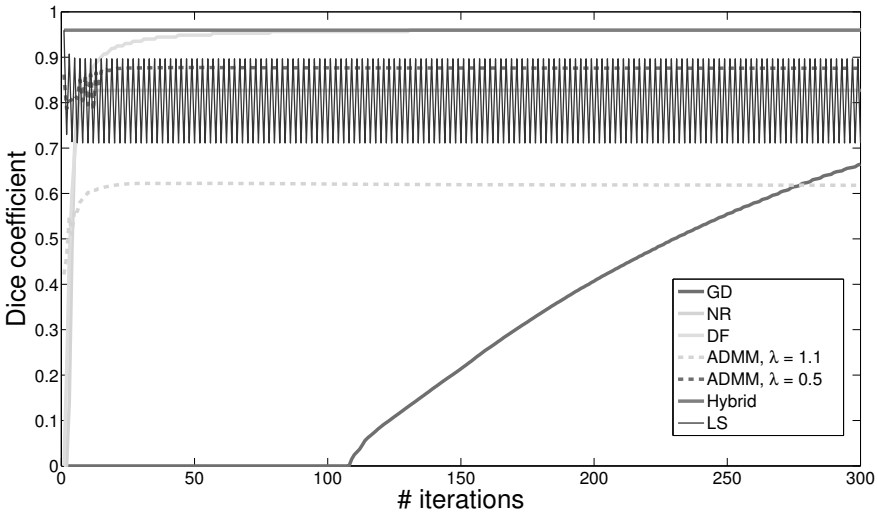


Figure 5.15: A comparison of the convergence of the active contours for noisy images with stronger regularisation. The Dice coefficient between ground truth and the active contours are plotted for each iterations step. The methods used for the active contour optimisation are: gradient descent, Newton-Raphson, line search, dual formulation, ADMM and a hybrid method combining line search and Newton-Raphson.

optimisation and the hybrid method combining line search and Newton-Raphson result in the best convergence. The convergence of the ADMM method clearly varies depending on the optimisation parameter λ. Note that both λ values resulted in reasonable convergence in subsection 5.3.7.2, while $\lambda = 1.1$ results in poor convergence on noisy images. So in order to have good convergence, the ADMM optimisation strongly depends on proper parameter tuning for each different application or image dataset.

5.4 Conclusion

This chapter discussed image segmentation using geometric active contours. It describes the background on implicit contour representation, e.g. using level-sets and characteristic functions, and how these representations can lead to active contours with convex energy terms. We proposed a new probabilistic interpretation to the active contour framework, which represent active contour segmentation as an inverse problem. In contrast to other geometric active contour formulations, the proposed method relies on probability theory, instead of ad hoc weight settings to tune the segmentation method for a specific application. We show that many state-of-the-art segmentation methods can be derived from this framework by assuming

certain probability distributions, e.g. the active contours without edges [Chan 06], smooth Mumford-Shah segmentation [Bresson 07b], texture based geometric active contours [Paragios 02], dictionary based active contours [Xia 12], statistical shape prior segmentation [Pock 10], etc. The proposed probabilistic framework justifies why for certain applications geometric active contours reported such accurate segmentation results. Furthermore we showed that the proposed model is more general than the conventional energy based active contours:

- The segmented objects do not require having the same foreground and background image models (see Fig. 5.4 for a clear example)
- New image models can also fit in this framework, e.g. polynomial image models, or local features such as generalised profiling [Babin 13b]. In contrast to [Deboeverie 13a], the proposed framework not only relies on the image model for accurate segmentation, but also incorporates prior information on the shape of the segment.

Finally we elaborately discussed several optimisation methods. For the first time, these methods where experimentally validated and compared for the segmentation of both regular as noisy images. A hybrid combination of heuristical optimisation and Newton-Raphson converges significantly faster than state-of-the-art optimisation methods, while taking less time to calculate a single iteration than most methods, e.g. the popular ADMM. This hybrid approach also needs significantly less memory than other state-of-the-art-methods, e.g. the dual formulation method.

The work discussed in this chapter is published in a number of international conference proceedings [De Vylder 10g, De Vylder 11a, De Vylder 11g, De Vylder 11b, De Vylder 12a, De Vylder 12e, De Vylder 14c], in a peer reviewed journal publication [De Vylder 13a] and in a journal publication in preparation [De Vylder 14a].

6

Detection and segmentation of cell nuclei

Chapters 4 and 5 mainly focused on generic segmentation techniques, often restricting the output to a limiting number of classes, e.g. foreground and background or a specific leave. For many biological applications, the number of objects in the image is unknown. Hence, different segmentation strategies are required. Such an application is the segmentation of cell nuclei. In this chapter we will discuss different strategies to segment nuclei in fluorescence microscopic images. We chose this application because the importance of cell nuclei in modern day research can hardly be overestimated [Skloot 10, Meijering 12].

The nucleus is an organelle in eukaryotic cells, which make up the tissues in animals. It contains most of the cell's genetic material in combination with a wide variety of proteins able to regulate gene expression. As such, it can be seen as the central hub of the cell, containing all information to form and maintain life. The genetic material and proteins are enclosed by the nuclear membrane, which separates the nucleus from the rest of the cell. The organization of DNA within the nucleus, allows for easy visualization by staining the DNA using a fluorescent protein or an organic dye. This has the benefit that the visualization is independent of cell type and invariant to specific gene mutations, in contrast to lamina based visualization. Given the vital function of cell nuclei in combination with their easy visualization, they are extensively studied objects for a wide range of applications:

Kinetic analysis of organisms By tracking cell nuclei over time, it is possible to study the evolution of organs, tumours and even full organisms over time. This can be useful for the study of dynamic processes, such as the development of an embryo [Keller 08, Keller 10], or for the analysis of growth under specific conditions, e.g. for the purpose of genetic engineering or for measuring the influence of specific stress factors on specific phenotypes [Chen 06, Quelhas 10b, Dhondt 12a]. Note that this analysis is on a cellular level, which holds much more information than monitoring using macroscopic modalities such as VIS imaging or ultra-sound.

Forensic investigations A cell nucleus is an important marker for offender identification, since the cell nucleus contains DNA. Specially for DNA identification in sexual assault cases, cell nuclei detection has proven its merit [Vandewoestyne 09a, Vandewoestyne 09b]. In these situations either a special type of cell nucleus, e.g. from a sperm cell, or a nucleus tagged using a special marker which only targets male cells has to be detected.

Study of nuclear pathologies The nuclear membrane is supported by a fibrillar network, the nuclear lamina. Genetic mutations influencing this nuclear lamina result in deviations in the nuclear architecture. These pathologies result in nuclei with aberrant shapes and appearance, which are related to diseases such as muscular dystrophies, lipid dystrophies, neuropathies and premature ageing syndromes [De Vos 10a]. Analysis of these pathological nuclei can result in a better understanding of these diseases or could serve for diagnostic purposes.

Pharmaceutical research In the case of pharmaceutics, automatic detection of specific cell types in the blood can shed light on the potential role of these cells in specific diseases [Merchant 02, Lepez 11]. An example of this is the amount of fetal microchimeric cells (i.e. cells that originate from a different individual) in the maternal circulation, which relate to thyroid disease, an autoimmune disease. However, it is very labour-intensive to locate one or more cells of a specific phenotype in a population of millions of regular cells. Staining the nuclei of the specific target phenotypes using fluorescent dyes can facilitate the detection of target cells.

This brief list is not exhaustive, but clearly illustrates the versatility in applications. Due to the biological importance of cell nuclei, several automated detection methods have been proposed in the past. Although the visualization is independent of the cell type, this does not mean that the intrinsic properties do not differ between cell types. As a consequence, different cell types result in different appearance and organization of cell nuclei. Due to this variability, there exists no golden standard segmentation method, instead the optimal segmentation technique depends on the

application. Cell nuclei segmentation methods can generally be categorized into two groups: edge based and intensity based. The first group starts by detecting edges and fits a specific shape model to them [Li 08, Gladilin 08, Gudla 08, Quelhas 10a, Tabus 10]. Both the use of binary edge detectors and edge strength maps have been proposed. The performance of these methods strongly depends on the quality of the edge maps, which is not always sufficient to detect individual nuclei in case of densely clustered nuclei. The second group first segments any fluorescent staining from the background. In contrast to isolated nuclei which can be detected using connected component labelling, detection of an individual nucleus in a cluster of touching nuclei requires an extra step. This is mainly done by requiring the detected segments of the cell nuclei to have a convex shape [Selinummi 05, Chen 06, Cloppet 08a, Cloppet 10, Kamentsky 11]. A different approach to split the nuclei clusters is by using prior knowledge from different channels, e.g. by segmenting cell walls or cytoplasm together with the assumption that each cell can only contain a single nucleus [Quelhas 10a, Pound 12]. However, this approach limits the application area since multiple stainings are necessary. Moreover, the methodology was only validated for specific cell networks.

For some applications it might not be necessary to analyse each individual nucleus, instead focusing on the statistics of a group of nuclei. In such a case, a confidence measure can be assigned to each detected nucleus, which then can be used to discard "unreliable" image segments and only analyse "reliable" nuclei detections. This however has the risk of rejecting specific sub-populations of cells or specific phenotypes [Ochoa Donoso 10, De Vos 10b]. So even in these situations, robust and accurate segmentation is still desirable.

As already mentioned, the appearance and organization of cell nuclei can be very different depending on the application, thus requiring different segmentation algorithms. In this chapter we will discus several segmentation strategies starting from different use cases. In the next section we will discuss two use cases where edges are the prominent feature used to segment individual cell nuclei. Section 6.2 discusses several methods focussing on intensity as salient feature. Three different use cases are discussed, showing common challenges such as non-uniform lighting, or the forming of dense clusters of nuclei both in 2D and 3D images. Section 6.3 concludes and gives a summary of the most important contributions discussed in this chapter.

6.1 Edge based nuclei detection

By dyeing the DNA or lamina using a fluorescent marker, a clear, but non uniform intensity is achieved. Since the density of both DNA and lamina is related

to the size of the nucleus, the intensity of the nuclei changes as the size of the nucleus changes during mitosis. To avoid this intensity variability, the use of more robust and stable features is preferred. Many different features have been proposed in state-of-the-art-methods, e.g. texture, concavity, etc., but specially the use of edges has been promoted in literature [Lin 03, Gudla 08, Yap 08, Sun 09, Hukkanen 10, Brullmann 12, Cattrysse 13]. These edge based methods consist of two independent steps: edge detection and fitting a model or template to the found edges. In this section we will describe two different approaches to edge based segmentation: the first assumes binary edge maps, i.e. a pixel belongs to an edge, or it doesn't, while the second approach calculates a measure of edge evidence for each pixel. Both approaches will be described on different use cases.

6.1.1 Use case: fibroblast nuclei

Fibroblasts are the most common cells of connective tissue in animals. They can be easily cultured and as such can be used to study the subcellular protein locations or the influence of external stress, such as radiation [Dieriks 09, De Vos 10b, Dieriks 10]. These studies can be related to cellular condition and by extension, to human health [Giuliano 09]. Normal, healthy fibroblast cell nuclei can be visualized by tagging the lamina using a fluorescent protein, resulting in good contrast images with clear edges. This advocates the use of edge based segmentation.

These methods typically start by detecting edges in the image. The most common approach to do this, is to apply a first-order difference operator as the estimation of the image gradient. Different linear filters can be used for this, e.g. Sobel, Prewit, Roberts' cross, etc., however these different linear operators do not result in spectacularly different edge detection results. Better results have been achieved using more complex edge indicators, e.g. based on local variance or multi-scale analysis [Lindeberg 98, De Vylder 10g]. However all these methods result in a similar output, i.e. a vector with all elements corresponding to the edge strength in a specific direction. The magnitude of this edge indicator vector is binarized using thresholding. Hysteresis thresholding has been proven useful in the envisaged application: it allows to remove small edges due to clutter and noise, while still resulting in long connected edges, even at locations where edges are slightly blurred [Canny 86]. Using morphological operators, the found edges are thinned to a line of one pixel width.

Using these edges, the location of individual nuclei borders is calculated. Note that edge segments can belong to multiple nuclei. Hence a shape model is required to detect individual nuclei. This shape model has to be fitted onto a part of the edges. This can either be done by intelligently scanning the parameter-space of the shape model, e.g. using a Hough-like transform, or by subsampling the edges, such that

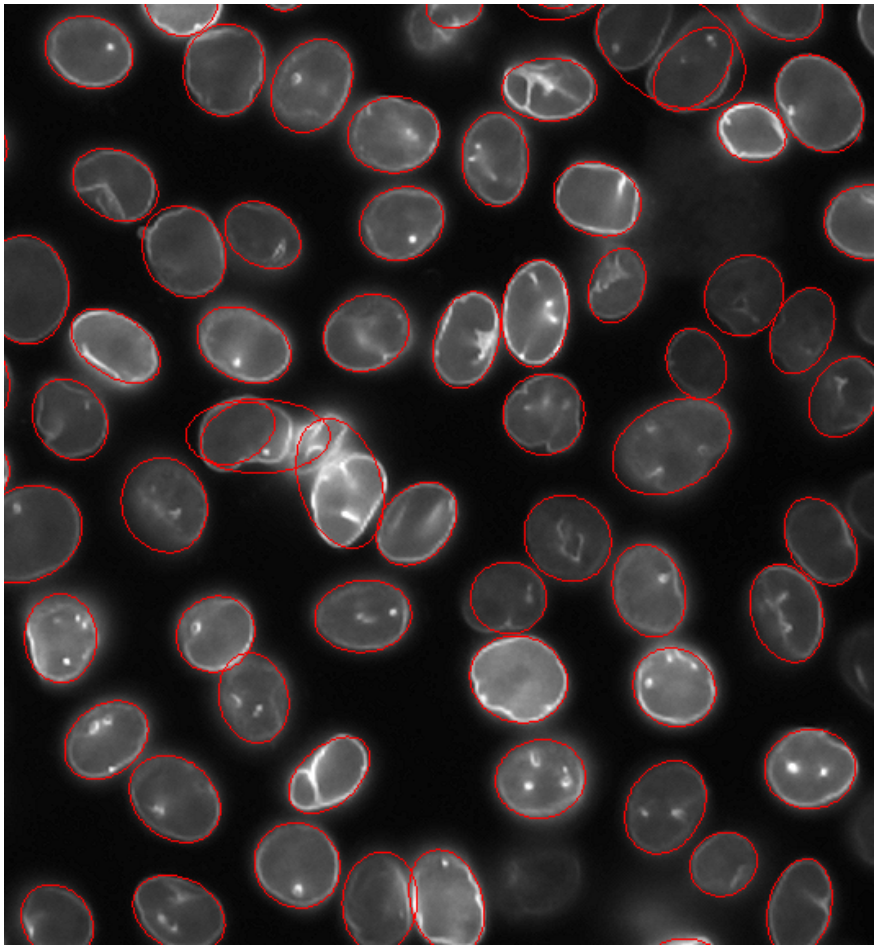


Figure 6.1: An example of a binary edge based segmentation result in combination with an elliptical shape model.

all samples presumably correspond to a single nucleus and then fitting the shape model to these samples. This second approach can be achieved using both an edges splitting strategy or by a RANSAC like strategy [Fischler 81, Hukkanen 10, Deboeverie 10].

The accuracy of these methods depends on a number of factors: the quality of the edge map, the shape prior; typically an elliptical nucleus is expected, which is sometimes too restrictive; the ability to subsample edge points from a single nucleus, etc. Even with these constraints, these methods have reported good results for specific applications. In Fig. 6.1 an example of a binary edge based segmentation is shown. As can be seen are most nuclei correctly detected. However in some locations the method gives erroneous results, this is mainly due to the use of edge segments that are part of multiple nuclei.

6.1.2 Use case: Hybridoma

Another interesting pharmaceutical application is the production of antibodies. Antibodies are a crucial part of the immune system. They have been proven useful both for diagnostic and for therapeutic purposes, e.g. in pregnancy tests [Chard 92], immunotherapy [Weiner 10], the diagnosis of tumours [Schrama 06] and auto immune diseases [Campbell 11], etc. These antibodies are typically produced using hybridoma technology: cells with the right antibodies, hybridomas in this situation, are selected and cloned. After culturing the hybridomas, the cells are tested for the secretion of the right antibodies. This procedure suffers from a number of drawbacks. Firstly, the screening and cloning are labour intensive and time consuming steps, resulting in a high production cost. A second drawback is that due to erroneous selection, cells lacking the correct immuno properties can survive the culturing, resulting in false positives in screening arrays. By carefully selecting only the hybridomas that produce the right antibodies, before culturing them, the process would be much more efficient both in time and in cost. This requires the screening of large populations of cells for the few cells with the right properties. This however is only feasible if the cells can be analysed automatically and preferably with a high throughput.

The previous proposed segmentation method is useful for images with clear, in focus edges, e.g. by tagging the lamina (the nuclear membrane). The segmentation accuracy depends on the ability to extract the edges. For nuclear studies based on DNA staining, the edges can be less clear due to blur or non uniform staining. By binarizing the edge evidence map, potentially interesting information is lost, e.g. the presence of an intensity difference between two pixels. To compensate for this information loss, a restrictive shape model is used to calculate the nuclear borders. In [Li 07b, Li 08] a segmentation approach was proposed that does not require

a binary edge map, but instead works on continuously valued edge maps, thus avoiding the drawbacks of binarisation. This method is inspired by the parametric active contour framework (see chapter 4), but rather than displacing contour points according to a force, a seed point for each pixel is displaced. All pixels whose seed point converges to the same end location are considered to belong to the same segment. A force based on the gradient vector flow (GVF) (see sec. 4.1.2 for more details) has shown to result in proper nuclei segmentation [Li 07b]. Instead of using the GVF-force itself, the force field is inverted. So instead of having a force towards edges, the force points away from edges. An example of this is shown in Fig. 6.2, where the force field is shown in red. Seeds moved under the influence of this force converge to stable locations, i.e. a locations where no force is active on the seed. All seeds moving towards the same stable location are part of the same segment. Such stable location either correspond to the centroid of a nucleus or to a region between two nuclei. Based on intensity or on the size of the segment, it is rather straightforward to distinguish if a stable location corresponds to a nucleus or to background. Such stable locations corresponding to a nucleus are called sinks, and are marked in green in the figure.

6.1.2.1 Edge propagation

Although this force field has been proven useful, it comes at a cost: the GVF force is defined as the solution of a partial deferential equation that is solved by iteratively optimizing of the complete force field. Solving this partial differential equation is both memory and time consuming. More computational efficient force fields have been proposed for parametric active contours, see section 4.5, however even with an efficient force field, this method is less suited for high-throughput segmentation. Aside from the calculation of the force, another computationally expensive step is required: all seeds have to be iteratively displaced according to the force field. In the remaining of this section, we propose a similar approach, but instead of using forces we will use an energy based on feature propagation (*fp*). This energy is computationally efficient to calculate and does not require an iterative updating of seeds.

To achieve this, we start from a feature function $f(x, y)$, e.g. the magnitude of the gradient of an image, where the feature map expresses how much *evidence* there is that there is an edge at pixel (x, y). The goal is to create an energy map, $g_{fp}(x, y)$, such that edge evidence is highest at the edge itself and gradually decreases further away from the edge. The main idea is to propagate strong edge evidence at a certain pixel to its neighbouring pixels with lower edge evidence. This step propagates edge evidence to direct neighbours. This can be efficiently done using the following dual scan line algorithm:

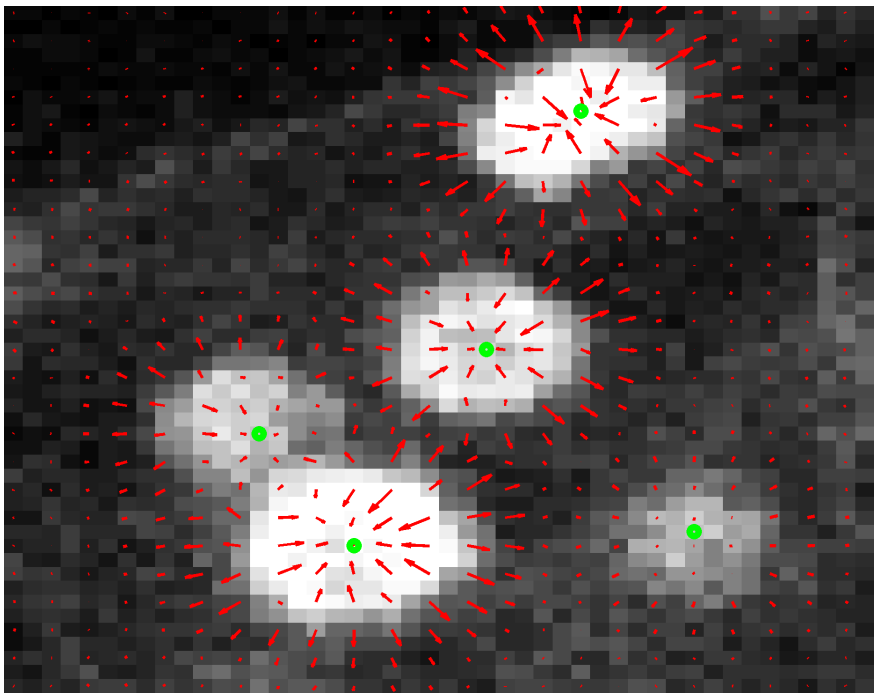
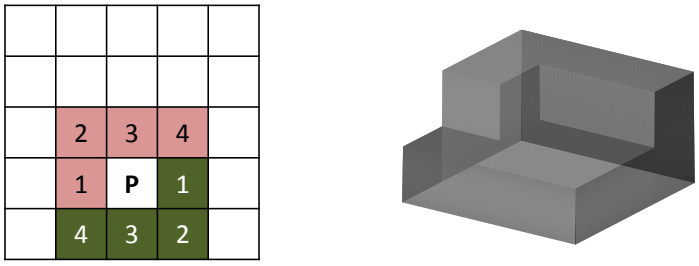


Figure 6.2: An example of the inverse gradient vector flow. The force points outwards the nuclear boundaries. Points where the force is zero, i.e. where opposing forces neutralize each other are shown in green. These so called sinks typically correspond to the centroid of an individual nucleus.



(a) (b)

Figure 6.3: (a) An example of processed neighbourhoods used by 2D feature propagation. The red one is used in the first run, the green neighbourhood is used for the second run. (b) The first run neighbourhood for 3D segmentation [De Vylder 11e].

1. In order to initialise our scanning algorithm, set the energy values for the first row and for pixel $(2, 1)$ equal to zero, i.e. $g_{fp}(1, y) = 0$ and $g_{fp}(2, 1) = 0$.
2. Scan the feature map row by row from top to bottom
3. In each row, the pixels are scanned from left to right
4. We first define the neighbourhood function $n(x, y, i)$, which returns neighbouring pixel locations that are already processed according to the scan line order, i.e.

$$n(x, y, i) = \begin{cases} (x, y-1) & \text{if } i = 1 \\ (x-1, y-1) & \text{if } i = 2 \\ (x-1, y) & \text{if } i = 3 \\ (x-1, y+1) & \text{if } i = 4 \end{cases}. \tag{6.1.1}$$

Fig. 6.3(a) shows the already processed neighbour pixels for $\mathbf{p}$ in pink. Using this neighbourhood function, we can calculate the energy value for pixel (x, y) as follows:

$$g_{fp}(x, y) = \max \begin{cases} \max_i \gamma \, g_{fp}(n(x, y, i)) \\ f(x, y), \end{cases} \tag{6.1.2}$$

where $\gamma \in [0, 1]$ is a weighting coefficient, which determines the speed at which feature evidence decreases. Note that all pixels of $g_{fp}(n(x, y, i))$ are already processed in previous steps, due to the scanning order. This energy map propagates feature evidence of a pixel beyond its direct neighbours in the scanning direction.

This algorithm propagates feature evidence from top to bottom and from left to right of the image. Repeating the algorithm in the opposite scanning direction,

propagates the feature evidence in the remaining directions. In this second run, the neighbourhood function is defined as

$$n'(x,y,i) = \begin{cases} (x, y+1) & \text{if } i = 1 \\ (x+1, y+1) & \text{if } i = 2 \\ (x+1, y) & \text{if } i = 3 \\ (x+1, y-1) & \text{if } i = 4 \end{cases}. \tag{6.1.3}$$

This new neighbourhood is shown in dark green in Fig. 6.3(a). The proposed energy map has the property that feature evidence gradually decreases. This has the advantage that if there is a small place where the evidence is less pronounced, e.g. due to blur, it is strengthened by its neighbouring evidence. The proposed energy map is robust against noise since the propagated edge evidence is generally more pronounced than the weak gradients due to noise.

Note that this technique of evidence propagation is easily extended to 3D. The already processed neighbourhood would then involve the previous processed slice as well, i.e.

$$n(x,y,z,i) = \begin{cases} (x, y-1, z) & \text{if } i = 1 \\ (x-1, y-1, z) & \text{if } i = 2 \\ (x-1, y, z) & \text{if } i = 3 \\ (x-1, y+1, z) & \text{if } i = 4 \\ (x-1, y-1, z-1) & \text{if } i = 5 \\ (x-1, y, z-1) & \text{if } i = 6 \\ (x-1, y+1, z-1) & \text{if } i = 7 \\ (x, y-1, z-1) & \text{if } i = 8 \\ (x, y, z-1) & \text{if } i = 9 \\ (x, y+1, z-1) & \text{if } i = 10 \\ (x+1, y-1, z-1) & \text{if } i = 11 \\ (x+1, y, z-1) & \text{if } i = 12 \\ (x+1, y+1, z-1) & \text{if } i = 13 \end{cases} \tag{6.1.4}$$

$$n'(x,y,z,i) = \begin{cases} (x, y+1, z) & \text{if } i = 1 \\ (x+1, y+1, z) & \text{if } i = 2 \\ (x+1, y, z) & \text{if } i = 3 \\ (x+1, y-1, z) & \text{if } i = 4 \\ (x-1, y-1, z+1) & \text{if } i = 5 \\ (x-1, y, z+1) & \text{if } i = 6 \\ (x-1, y+1, z+1) & \text{if } i = 7 \\ (x, y-1, z+1) & \text{if } i = 8 \\ (x, y, z+1) & \text{if } i = 9 \\ (x, y+1, z+1) & \text{if } i = 10 \\ (x+1, y-1, z+1) & \text{if } i = 11 \\ (x+1, y, z+1) & \text{if } i = 12 \\ (x+1, y+1, z+1) & \text{if } i = 13 \end{cases} \tag{6.1.5}$$

An illustration of such a 3D neighbourhood is shown in Fig.6.3(b). It might be necessary to adjust the algorithm in order to cope with the significant difference between the resolution in the focal plane and the resolution in the z-axis. This can easily be achieved by using different weighting parameters γ_i in eq. (6.1.2), depending if the neighbouring pixel i is in the same plane or in the previous processed plane.

6.1.2.2 Nuclei segmentation

The original continuous edge based segmentation methods use a force field instead of an energy field [Li 07b, Li 08]. While edge based forces with better computational efficiency than the GVF force have been studied, see section 4.1.2, the actual segmentation, i.e. displacing seed-points until a stable location is reached, remains a computationally demanding task. The proposed energy field does not suffer this burden, since it can easily be segmented directly using the watershed algorithm, which is discussed in 3.2.3. Computationally efficient watershed algorithms have been reported in [De Smet 00, De Bock 07], these are efficient both in the sense of processing time and in the sense of memory requirements.

Although the proposed feature propagation method is robust to noise there is still a risk of over-segmentation, mainly with elongated elliptical nuclei. This can however easily be solved using region merging. For each segment it is calculated how much it would cost to merge the segment with a neighbouring segment. As the merging cost criterion, we propose the minimum value of the energy field along the border between both segments. Neighbouring segments with a merging cost below a certain threshold are merged. Both the calculation of the merging costs

between segments as the merging itself can be done efficiently using the data-structures proposed in [De Bock 06b].

This segmentation approach partitions the image in segments. The average segment intensity is compared with a threshold in order to discriminate between nuclei segments and background segments. For the selection of an optimal threshold value, we refer to [Sezgin 04].

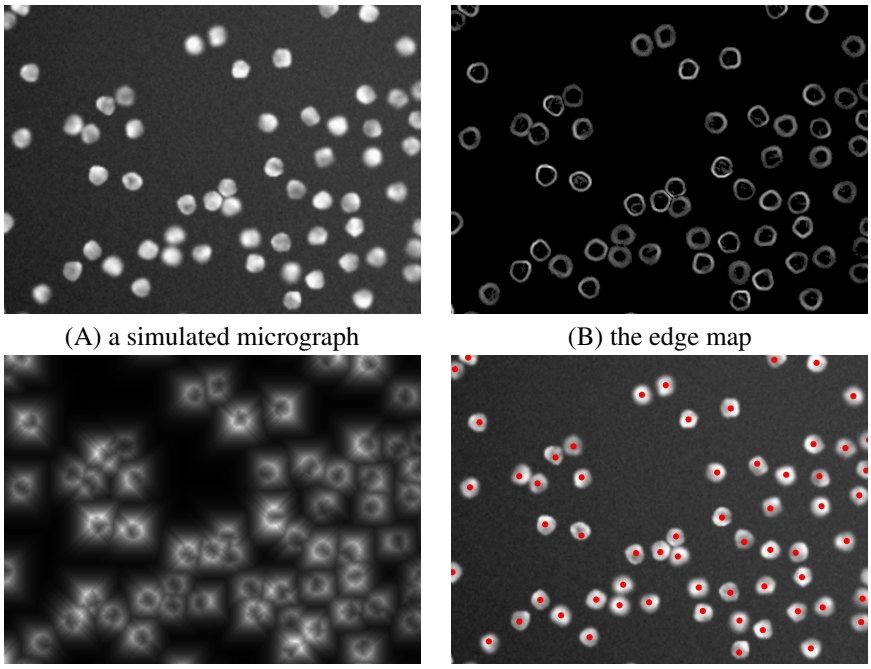
6.1.2.3 Experimental results

As a first validation of the proposed method, a synthetic data set is analysed [Ruusuvuori 08]. These synthetic images were proposed as a common benchmark for nuclei segmentation and detection. The synthetic images show the same intrinsic properties as real microscopic images of cell nuclei: blurred nuclei, non uniform intensity within a nucleus, touching nuclei, non uniform background, etc. Fig. 6.4(A) shows an example of such a synthetic raw image. Fig.6.4(B) shows the edge map of the simulated micrograph. Fig. 6.4(C) shows the energy map, i.e. the propagated edge map. Fig. 6.4(D) shows the centroids of the found segments, i.e. the detected nuclei. As can be seen, both isolated and clustered nuclei are correctly detected.

A second experiment was done using real 2D microscopic images of hybridomas (see appendix A). These images have a resolution of 1344×1024 pixels and contain on average 153 cell nuclei per image. Ground truth was acquired by us by manually annotating the nuclei.

The results of nuclei detection in both synthetic images as in real images are summarized in Table 6.1. The first two rows display the ground truth, i.e. respectively the number of images in the data set and the number of nuclei shown in the data set. The next two rows list the number of nuclei detected by the proposed algorithm, i.e. respectively the number of nuclei counted in the full data sets and the maximum difference between the amount of nuclei visible in the image and the amount of nuclei found in the image. The last two rows show the precision and recall of the full data set (see appendix B).

The proposed method is inspired by the method defined in [Li 08]. This method in [Li 08] consists of two steps: calculating a force field and then segmenting the image based on this force field. The use of a GVF force was proposed and validated. The calculation of the GVF field requires the optimization of a partial differential equation, which needs $O(N^2)$ operations for a single iteration for a square image with dimension $N \times N$. The number of required iterations depends on the nucleus size. Assuming that the maximum diameter of a nucleus is defined as $2M$, M iterations are generally needed to calculate a useful force field [Li 07a].



(A) a simulated micrograph (B) the edge map

(C) the propagated energy map (D) cell nuclei detection

Figure 6.4: An example of a simulated micrograph and the different steps used for the detection of the individual nuclei [De Vylder 11e].

	simulation	hybridomas
# images	20	4
# nuclei	6000	612
total count	6006	643
max error	3	11
precision	0.996	0.933
recall	0.996	0.982

Table 6.1: Nuclei detection results for both simulated data and real 2D microscopic images.

So the GVF-approach has an $O(MN^2)$ complexity. Other, computational less demanding force fields have been proposed in literature, e.g. vector field convolution (VFC) [Li 07a, Wang 08]. These convolution based force fields have a complexity of $O(N^2 \log N^2)$, which is determined by the complexity of the 2D FFT and IFFT algorithms used. Our proposed scanning algorithm recalculates the value of each pixel twice, resulting in an $O(N^2)$ algorithm. The proposed method has not only a low processing time, but also uses less memory. Only the current row/slice and the previous row/slice need to be stored in memory.

After calculating a force or energy, the nuclei still need to be segmented. This is generally done by following the force field towards sinks [Li 07b, Li 08]. This has to be done for each pixel. In the worst case a seed has to be displaced M times, i.e. the maximum radius from a nucleus. This results in an algorithmic complexity of $O(MN^2)$. In this section we have proposed to use an energy field instead of a force field. This simplifies the detection, since nuclei correspond with local minima in the energy field. These can easily be found using a watershed algorithm. There exist memory and computational efficient implementations of the watershed algorithm, resulting in a computational complexity of $O(N^2)$ [De Bock 05, De Bock 07]. The proposed method is linear in complexity, and is as such a good candidate for high throughput analysis.

6.2 Intensity based nuclei detection

To overcome the limitations of edge based segmentation, numerous methods have favoured intensity as a discriminant feature instead of edges [Carpenter 06,Chen 06, Lamprecht 07, Rohde 08, Yan 08, Keller 08, Cloppet 08b, Gudla 08, Meyer 09, Harder 09, Al-Kofahi 10, Nemeth 11, Pound 12, Hodneland 13]. While intensity is a common feature used by many different methods, these methods differ either in the way that they exploit the intensity, or in the way that they distinguish individual nuclei. An optimal segmentation method should take into account all intrinsic sample properties, such as the appearance, shape, the configuration of the nuclei, etc. These properties depend on the cell culture and on the acquisition method, e.g. in vivo, in situ or in vitro. Within this section we will elaborate on three different segmentation methods starting from specific use cases where nuclei analysis is important. The first method is developed to cope with varying intensity of nuclei, the second method will tackle dense nuclei clusters in 3D data sets and the final method will handle the problem of clustered nuclei in 2D images.

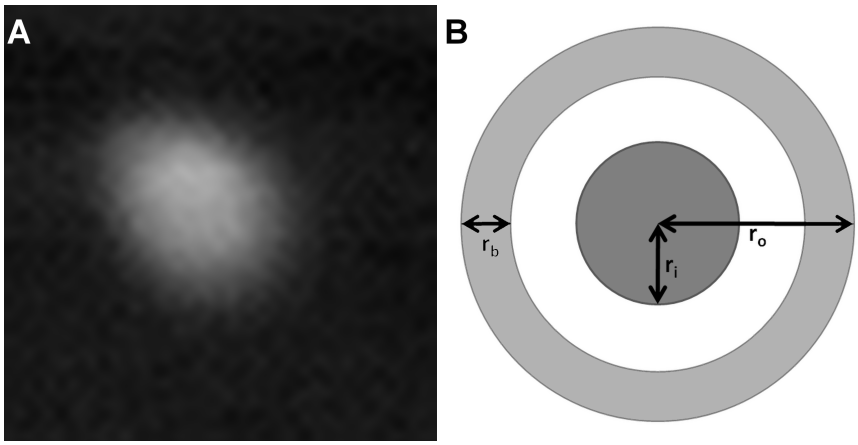


Figure 6.5: On the left a close up from a single nucleus is shown, the right depicts the linear filter used to normalize the image.

6.2.1 Use case: embryonic cell nuclei

The study of early organ development and embryogenesis is specially interesting at a cellular level. Such a cellular analysis comprises cellular displacements, cell division, morphogenetic analysis, local activity, etc. Both single organs and small models systems such as C. elegans have been reconstructed based on nuclei tracking [Lemaire 06, Rembold 06, Megason 07]. However the reconstruction of more complex vertebrate embryos poses a number of technical challenges. One of the major challenges is that analysis algorithms have to scale to large data: C. elegans embryogenesis requires the analysis of less than 700 nuclei, while an 18h old zebrafish embryo contains approximately 16000 nuclei. Due to the high nuclear mobility, a high temporal resolution is needed as well, i.e. up to one 3D image every 90s. Finally the central yolk forms an extra challenge. This yolk is opaque at the wavelengths used for both illumination and emission in fluorescence microscopy, thus partially obstructs the visualization of cell nuclei [Keller 08, Keller 10].

6.2.1.1 Local intensity

The nucleus intensity varies locally due to partial obstruction of light by the central yolk cell, as can be seen in Fig. 6.6.A. This intensity variability severely hampers intensity based segmentation methods. To overcome this challenge we propose to transform the intensity of a pixel based on the intensity of a predefined neighbourhood of that pixel. In concreto we attempt to take the difference between the intensity at a pixel and the average intensity of the background in the vicinity of

the pixel. To make the method more robust against noise, we propose to not use the intensity at the pixel, but the average of a small neighbourhood around the pixel. This result in the following transform:

$$g(\mathbf{x}) = \frac{1}{\#\Omega_{\mathbf{x}}} \sum_{\mathbf{y} \in \Omega_{\mathbf{x}}} f(\mathbf{y}) - \frac{1}{\#\Theta_{\mathbf{x}}} \sum_{\mathbf{z} \in \Theta_{\mathbf{x}}} f(\mathbf{z}), \tag{6.2.1}$$

where $\#$ represents the cardinality of a set, $\Omega_{\mathbf{x}}$ corresponds to the small neighbourhood around $\mathbf{x}$, and $\Theta_{\mathbf{x}}$ represents a set of background pixels in the vicinity of $\mathbf{x}$. In Fig. 6.5.B an example of such a neighbourhoods is shown. In dark gray, the foreground neighbourhood, corresponding to $\Omega_{\mathbf{x}}$, is shown, whereas the background pixels, $\Theta_{\mathbf{x}}$, correspond to a ring surrounding a nucleus, shown in light gray. The foreground region is chosen in such a manner, that when $\mathbf{x}$ corresponds to the centroid of a nucleus, that all pixels in $\Omega_{\mathbf{x}}$ belong to the nucleus as well, while all pixels belonging to $\Theta_{\mathbf{x}}$ are certainly not part of the same nucleus. An example of such sets can be defined as e.g.

$$\begin{aligned} \Omega_{\mathbf{x}} &= \{\mathbf{y} \mid \mathbf{y} \in \mathbb{R}^2, \|\mathbf{x} - \mathbf{y}\|_2 \leq r_i\} \\ \Theta_{\mathbf{x}} &= \{\mathbf{z} \mid \mathbf{z} \in \mathbb{R}^2, \|\mathbf{x} - \mathbf{z}\|_2 \leq r_o \wedge \|\mathbf{x} - \mathbf{z}\|_2 > r_o - r_b\}, \end{aligned}$$

where $\wedge$ represents the "and" operator, and r_b, r_i , r_o are defined according to Fig. 6.5.B. The transform proposed in eq. (6.2.1) can be implemented as a linear filter, thus resulting in a very efficient computation. This allows for high throughput, which is necessary given the high spatial and temporal resolution. The nuclei of the digital embryo data set are sufficiently isolated, i.e. they do not optically touch each other. Therefore computational efficient and simple segmentation methods such as thresholding (see chapter 3) can be used to segment individual nuclei.

6.2.1.2 Experimental results

The proposed method is validated using a data set monitoring the development of a zebrafish embryo. The data set is acquired using scanned light sheet microscopy, more detailed information on the data set can be found in appendix A. An example of a slice of the digital embryo is shown on the left of Fig. 6.6. As can be seen, not all cell nuclei are equally visible, this is due to the obstruction of illumination by the central yolk. The proposed method is robust against inhomogeneous lighting conditions, as can be seen on the right of Fig. 6.6, where the segmentation result is shown.

To validate the methods quantitatively, the digital embryo data set was manually annotated by an expert. This manual annotation resulted in the 3D location of the

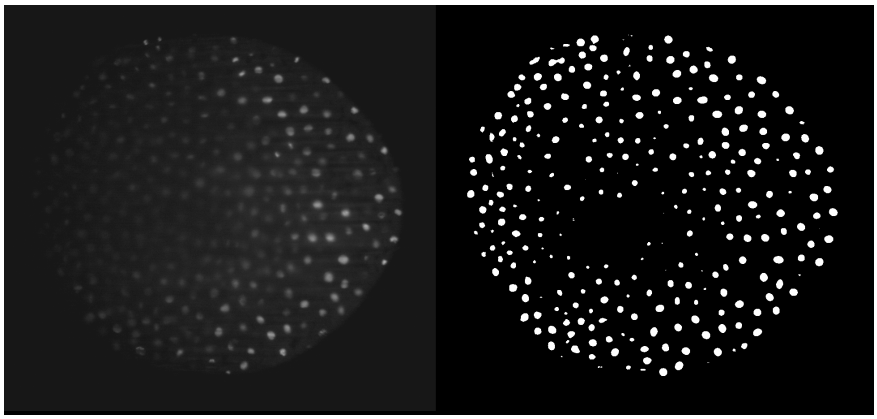


Figure 6.6: An example of a slice from a 3D micrograph of a zebrafish embryo. On the left, the cell nuclei visualized using DAPI stained fluorescence microscopy, on the right the segmentation result.

centroids of 1493 nuclei. Since touching nuclei are not an issue in this embryonic data, we quantitatively compare our method with the following isolated nuclei segmentation methods quantitatively: trainable Weka Segmentation [Bradhurst 08, Hall 09], maxEntropy [Kapur 85, Sezgin 04, Gudla 08] and moments preserving thresholding [Tsai 85, Sezgin 04]. The nuclei found by these methods, generally correspond to real nuclei, which results in good precision measurements, as can be seen in Table 6.2. Specially the maxEntropy scores good on this aspect. However, all these methods miss many nuclei due to the low intensity of some of the nuclei. This results in poor recall measures. The proposed method is more robust for the non-uniform lighting, resulting in better recall measures, while still achieving good precision. This combination results in a significantly better F-measure, as can be seen in Table 6.2.

Due to high dynamics of nuclei in embryonic development, the data sets not only have a high spatial resolution but also high temporal resolution. This results in huge amounts of data requiring high-throughput analysis methods. In order to measure the computational efficiency, all methods were run on an Intel i7 Q720 1.6 GHz CPU with 4GB RAM.[1] The processing time for a 117 megavoxel image is shown in the right column of Table 6.2. While the threshold based methods from literature are computational the most efficient, the proposed method is still reasonably fast, specially compared with more complex methods, such as the Weka segmentation algorithm [Bradhurst 08, Hall 09]. The proposed method is more than efficient enough to achieve "real-time" processing, i.e. a frame rate of a 3D microscopic image every 90s [Keller 08].

[1] All methods are implemented in Matlab R2007b1.

	count	P	R	F	speed (s)
trainable Weka Segmentation	953	0.843	0.538	0.657	3002.392
maxEntropy	958	**0.992**	0.636	0.775	**19.267**
moments based thresholding	1146	0.909	0.638	0.790	12.730
proposed	1342	0.920	**0.827**	**0.871**	28.026

Table 6.2: A comparison of different cell nuclei segmentation methods on the digital embryo data set (P = precision, R = recall, F = F-measure).

6.2.2 Use case: Neuronal cell nuclei

The previous use case studies temporal dynamics based on the motion of nuclei. However using different bio-imaging techniques, it is also possible to study temporal functional dynamics. Using new probes, it is not only possible to visualize nuclei, but also to distinguish the nuclei of excitatory neurons [Kim 09, Kwok 09]. These probes permit visualization of the recent history of activity in the brain at cellular resolution, thus allowing identification of not only which areas of the brain are activated during cognitive processing, but which specific neurons. This new functional neuroimaging modality provides a window on the physiological basis of cognition. In particular it can shed light on the mechanism of memory and memory disorders such as Alzheimer's disease. More specifically it can help us to understand the dynamic interactions among neuronal populations and synaptic plasticity mechanisms that underlie these phenomena.

In order to analyse cellular activity, the neuronal nuclei have to be segmented from the images acquired with these new imaging modalities. While these nuclei have a similar appearance as nuclei from Fibroblasts or Hybridomas (see section 6.1), their organization relative to each other differs significantly from in vitro observations. There are places in the brain such as CA1, CA3 and the dentate gyrus in the hippocampal formation where nuclei are densely packed, which typically results in erroneous segmentation results. Hence this requires a more specialized segmentation approach, specially dedicated for detecting nuclei in dense clusters. An example of such a dense cluster of neuronal nuclei can be seen in Fig. 6.7. This micrograph shows a maximum projection of a small section of a rat brain. Nuclei are labelled using DAPI (shown in blue) and the green channel corresponds to NeuN immuno, a neuron specific protein.

Due to the goal of activity mapping in the brain, it is necessary to monitor large areas of the brain, thus resulting in large data sets, i.e. slices of 13010×1024 pixels. To keep the analysis of this big data problem computational feasible, we will analyse the data on a slice-by-slice basis,only combining consecutive slices in a final phase.

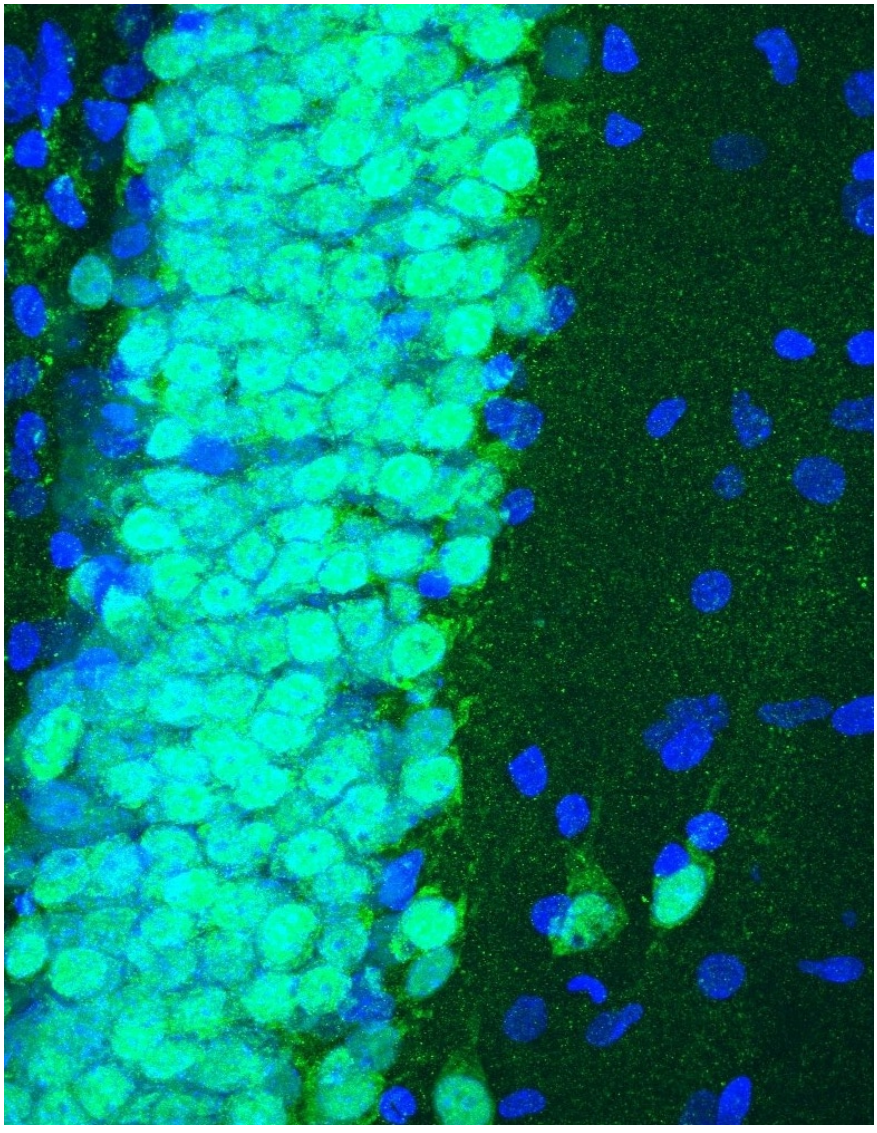


Figure 6.7: A maximum projection of a dense cluster of cell nuclei in a cross section of the dentate gyrus of a rat brain.

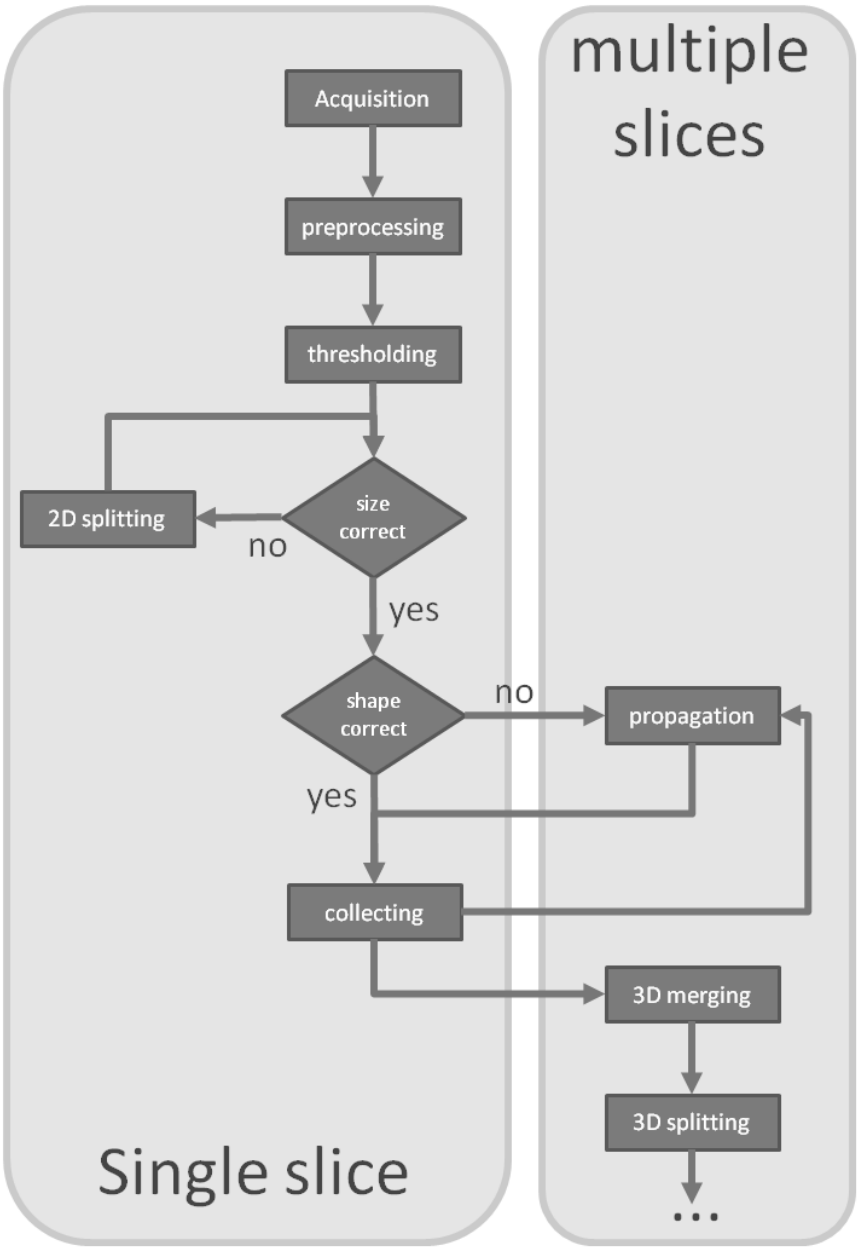


Figure 6.8: An overview of the 3D segmentation work flow.

In contrast to methods from literature tackling large groups of (clustered) nuclei, we do not put constraints on the possible spatial configuration of the cells [Quelhas 10a, Meijering 12, Pound 12]. Since these assumptions are not always vallid, we do not assume that the cells are organised in a regular grid, or that within a known region there should be exactly one nucleus. Fig. 6.8 shows the overall work flow of the proposed method. First the slices are individually preprocessed in order to improve the contrast between nearly touching nuclei. The preprocessed slices are then segmented using a simple thresholding approach. The resulting segments are then checked: segments which are to big are considered as potential clustered nuclei and are passed to a splitting step. After a number of iterations of checking and splitting, all segments have an area which fits within the expected range. Then all segments are checked a second time, but this time based on a number of shape features. Segments which are considered likely erroneous (whose shape features deviate from what is expected) are discarded and replaced by the contours of segments in neighbouring slices located in the same place. This initial contour is then updated such that the contour delineates the nucleus in the current slice. Note that this is the first step incorporating information from other slices. Then the propagated contours and the segments that passed the check are passed to a merging step, which combines 2D segments of different slices into full 3D segments. To correct for potential false merges, an extra 3D splitting step concludes the work flow. In the next part of this chapter we will elaborate in more detail on the different steps mentioned in this work flow.

6.2.2.1 Contrast enhancement

To enable accurate segmentation, we propose to enhance the contrast between the background and the nuclei. Especially at locations where nuclei nearly touch each other, this step will improve the segmentation. A variety of techniques have been proposed to improve the contrast in an image, e.g. gamma-correction, homomorphic filtering, histogram equalization, histogram mapping, retinex, etc. [Jobson 96, Jobson 97, Murtagh 03, De Vylder 10g]. However, these images pose some extra challenges: the contrast enhancement step should be able to cope with differences in DNA-density, i.e. differences in nuclei intensity, differences in focal depth, potential bleed through, etc. In concreto it would be desirable to enhance the contrast between nearly touching cell nuclei but not inside a nucleus. The regions between two touching nuclei typically correspond to a dark region with only dark neighbouring pixels in one direction and in the opposite direction the dark region is surrounded by brighter regions. An example of this is shown in Fig. 6.9, where the intensity profiles of the red and green lines are plotted on the right. The bottom profile, i.e. the red line, corresponds to a relative flat region, with low values. The green line (the top intensity profile) corresponds to a small region of

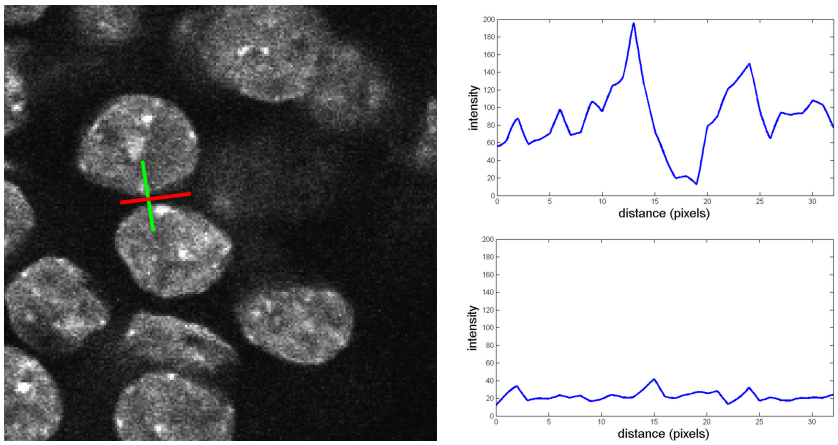


Figure 6.9: An example of the local structure of densely clustered nuclei in a fluorescence micrograph. An the left, an example of a DAPI stained image. On the top right the intensity profile corresponding to the green line in micrograph, while the bottom right plot corresponds with the the intensity profile of the red line.

low intensities, flanked by regions of higher intensity. This type of local structure corresponds to a ridge, for which many automatic detectors have been proposed in literature [Koenderink 94, Eberly 94, Steger 98, Frangi 98]. A common approach is to calculate a "ridge-measure" based on the local geometric structure in the image. This local structure is well described using the Hessian operator. The Hessian of a pixel, (x, y), of an image, $f(.,.)$, is defined as

$$(\mathcal{H}f)(x,y) = \begin{bmatrix} \frac{\partial^2 f(x,y)}{\partial x^2} & \frac{\partial^2 f(x,y)}{\partial x \partial y} \\ \frac{\partial^2 f(x,y)}{\partial x \partial y} & \frac{\partial^2 f(x,y)}{\partial y^2} \end{bmatrix}. \tag{6.2.2}$$

By analysing the eigenvalues of the Hessian, the principal directions of the local second order structure can be decomposed [Frangi 98, Lindeberg 98]. Let λ_1 and λ_2 denote the two eigenvalues with $|\lambda_1| \geq |\lambda_2|$. The relation between the eigenvalues and the local structure is shown in Table 6.3.

Several different ridge-measures based on the Hessian's eigenvalues have been proposed in literature, e.g. main principle curvature [Koenderink 94, Eberly 94, Steger 98, Frangi 98]. For our purpose we propose the Frangi vesselness measure [Frangi 98]. This measure combines two measures, i.e. a geometric structure and an intensity based measure. In order to measure the geometric structure the following blobness measure was proposed

$$m_g = \frac{|\lambda_2|}{|\lambda_1|}. \tag{6.2.3}$$

λ_1	λ_2	local structure
N	N	noisy, no preferred structure
H^-	L	ridge structure, bright
H^+	L	ridge structure, dark
H^-	H^-	blob structure, bright
H^+	H^+	blob structure, dark

Table 6.3: Possible local structures in an image, where $+/-$ indicates the sign of the eigenvalue and where H= high, L= low, N= noisy, typically small.

This measure has a high value if the pixel corresponds to a blob like structure, and has a value close to zero for ridge like structures. Note that this value is invariant to intensity rescaling. Hence, this filter will also have strong filter response in the background, due to noise and clutter [Frangi 98]. Since the magnitude of the second order derivatives is typically small in the background, the Frobenius matrix norm can be used as a measure of belonging to the foreground, i.e.

$$m_i = \|\mathcal{H}\|_F = \sqrt{\lambda_1^2 + \lambda_2^2}. \tag{6.2.4}$$

This measure will be low in the background, due to the lack of structure. High contrast regions will result in at least one eigenvalue with large magnitude, and thus this measure will be high as well. Frangi et al. proposed to map these two measures, m_g and m_i, into probability-like estimates. The first measure, m_g is modelled using a Gaussian like distribution centred around zero. For the second measure, m_i, the likelihood that a pixel has low contrast is modelled using a similar Gaussian like distribution around zero. Since the measurement should actually highlight high contrast pixels, this Gaussian like measurement is subtracted from one. These new measurements are then combined into a single ridge measure:

$$m_r = e^{-\frac{m_g^2}{2\sigma_g^2}}(1 - e^{-\frac{m_i^2}{2\sigma_i^2}}). \tag{6.2.5}$$

Note that although inspired by probability theory, this measure is not the real likelihood of a pixel belonging to a ridge, e.g. the correct scaling factors are missing. This however does not seem to hamper its usability [Frangi 98, Frangi 99, Sofka 06].

We propose to estimate the likelihood that a pixel belongs to a nucleus, by combining this ridge measurement with the intensity using a simple linear regression model. In our regression model the intensity is directly proportional and the ridge feature is inversely proportional to the likelihood. While this model is a simplification of a proper trained regression model, it still results in images with good properties for segmentation. The contrast between background and nuclei is drastically

increased, while the nuclei are generally only effected by this preprocessing step at the location of nucleoli, which are typically encapsulated in the nucleus. An example of this is shown in Fig. 6.10, on the top the raw fluorescence micrograph is shown, in the middle image the output of eq. (6.2.5) is superimposed in red, and the bottom image shows the estimated likelihood image. As can be seen are the nuclei edges more crisp and are nuclei generally better separated by dark background.

6.2.2.2 Nuclei detection

A common approach in intensity based segmentation is to first apply an initial segmentation step where all nuclei pixels result in a single segment and all background pixels correspond to a second segment. In a second phase the nuclei segment is split into multiple segments, each corresponding to a single nucleus. This approach has the benefit of simplicity, while allowing adequate results for a number of applications [Carpenter 06, Cloppet 08b, Cloppet 10, Kamentsky 11, Meijering 12, Li 13].

Initial segmentation Since this step mainly focuses on distinguishing between two classes: foreground and background, simple and computational efficient segmentation methods are typically sufficient. Hence we propose the use of a Maximum Likelihood classifier:

$$\hat{c}(\mathbf{x}) = \underset{c(.)}{\arg\max}\, p(f(\mathbf{x}) \mid c(\mathbf{x})), \tag{6.2.6}$$

where $c(.)$ represents a characteristic function, resulting in zero if the voxel belongs to the background and in one if it corresponds to a nucleus. Since stained nuclei are generally brighter than the background, this can be efficiently implemented using thresholding, where the optimal threshold is estimated using a training data set. For the automatic estimation of a proper threshold in the absence of training data, we refer to the following works [Sezgin 04, Abramoff 04].

The preprocessing step sometimes finds ridges between bright blobs, corresponding with telomeres. As a consequence the preprocessing step can create dark regions within the nucleus. The maximum likelihood method classifies these darker regions typically as background. To compensate for this side effect of the preprocessing step, all holes in the segments are filled, e.g. using morphological reconstruction (see section 3.3.1.2).

Iterative splitting To split clustered nuclei into individual segments several approaches have been proposed. These techniques generally attempt to detect points

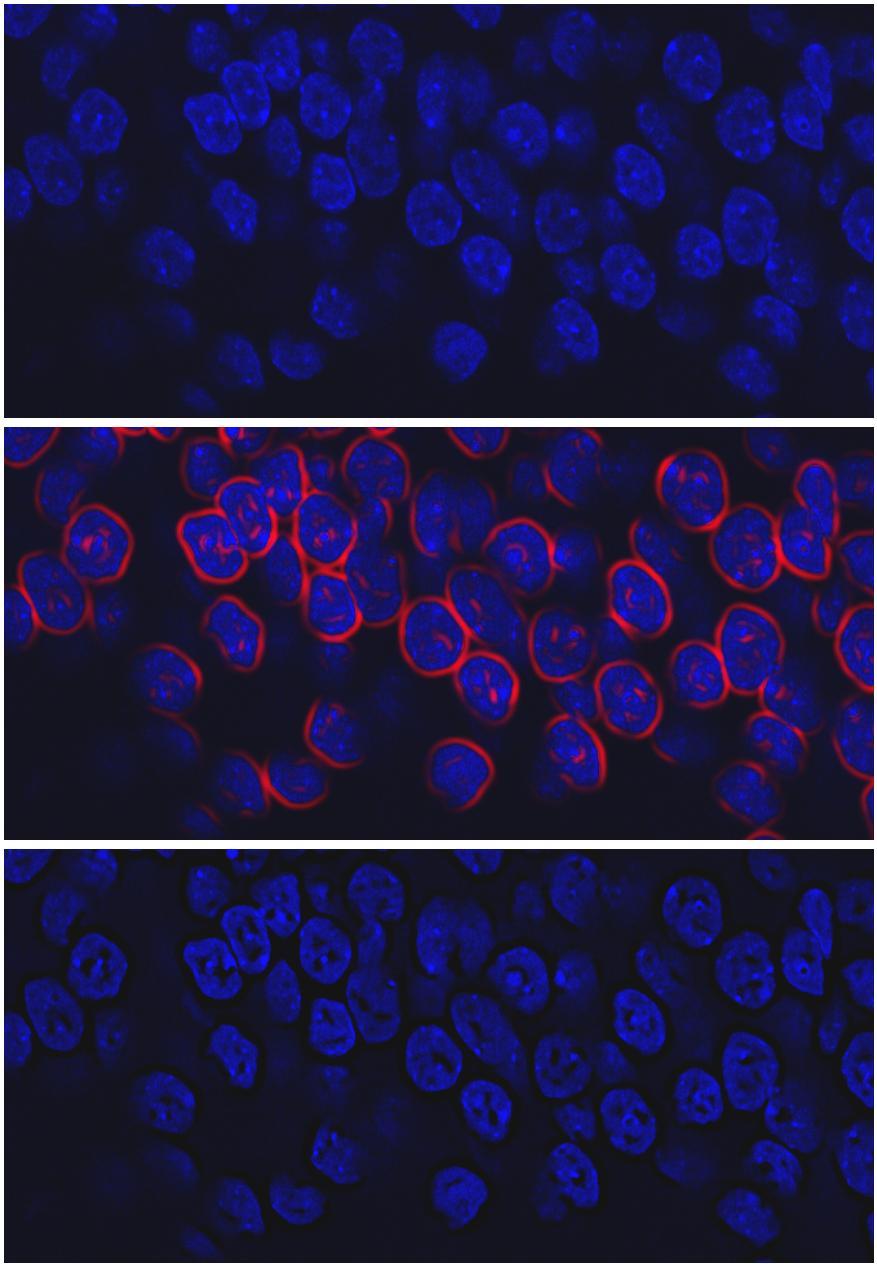


Figure 6.10: An example of the different preprocessing steps. The top image shows an example of a DAPI stained fluorescent micrograph, the middle image superimposes the output of the ridge filter in red and the bottom image represents the output of our regression model.

where the contours of two nuclei touch each other, called break points. By connecting the right break points the clustered nuclei can be separated correctly. Splitting techniques from literature can be grouped in two main categories: explicit and implicit methods.

With explicit methods the break points are first detected, typically by locating boundary points of high curvature [Cloppet 08b, Cloppet 10, Shu 13]. This however requires accurate boundary detection, a difficult task, specially for cell nuclei which are located at the border of the focal plane or for aberrant nuclei with deviating shapes. A second non trivial challenge is to correctly connect the correct break points, in order to split the segment into individual nuclei.

The implicit methods attempt to avoid these hazards by avoiding the accurate detection of break points, instead they assume that the correct break line corresponds to a line connecting concave parts of the segments boundary, in such a way that the connecting lines have the smallest possible length [Gudla 08, Meyer 09, De Vos 10b, Ollion 13]. By cleverly using the Euclidean distance transform and the watershed segmentation technique (see chapter 3), these methods are able to achieve their goals in a computational efficient way, without actually detecting the break points.

This is illustrated in Fig. 6.11. The likelihood image is shown in A, which is segmented according to the procedure discussed in the previous subsection, shown in B. Based on this binary image, a Euclidean distance map can be calculated, i.e. a map where the intensity of each pixel corresponds to the Euclidean distance to the closest background pixel. This is shown in Fig. 6.11.C. By applying the watershed segmentation method on the inverse of this distance map, the clustered nuclei are separated. However due to noise, clutter and deviations from the elliptical shape, this approach has the tendency to over-segment the nuclei, as can be seen in Fig. 6.11.D. To overcome this problem, the approach can be adjusted such that only break lines with a small length are considered [Abramoff 04, Meyer 09, De Vos 10b, Kamentsky 11]. The result of this length constrained approach is shown in Fig. 6.11.E. The result strongly depends on the maximal allowed length of a break line. The threshold on this length should be in relation to the morphology of the nuclei. For the case of brain nuclei this forms a challenge, since the brain consists of many different types of cells, and as a such, of nuclei with different types of morphology. An example of this is shown in Fig. 6.11.F. where a small elongated nucleus is depicted. The threshold on the length of the break line should be set such that even this nucleus is not over-segmented, however the width of this nucleus is in the same range as the length of the break lines in Fig. 6.11.B. So either these elongated nuclei are over-segmented, or the clustered nuclei are not separated.

To avoid this poor trade-off between over-segmentation and insufficient splitting,

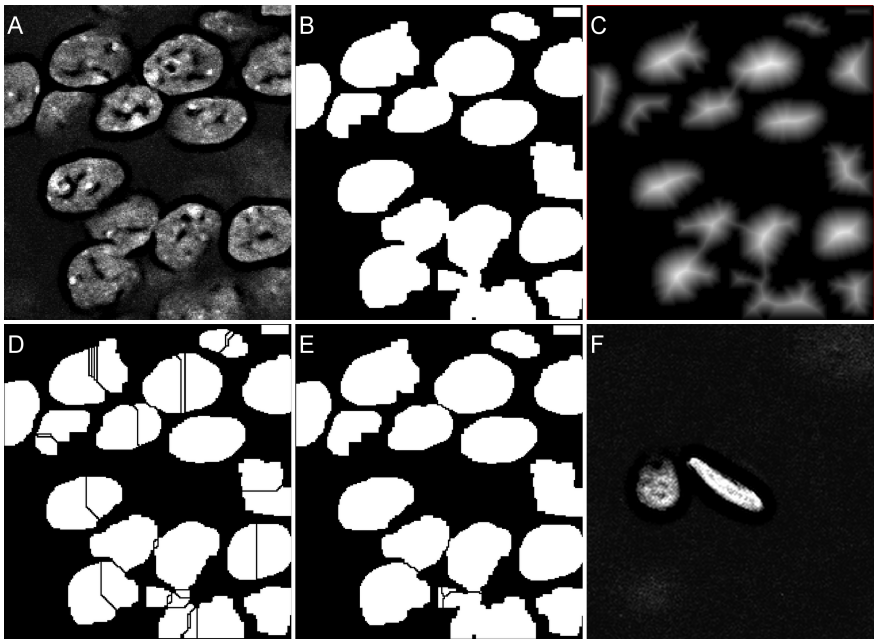


Figure 6.11: The different steps in individual nuclei segmentation: A) the nuclei likelihood map, B) the initial segmentation resulting in a binary image showing the foreground, C) the Euclidean distance map, D) watershed based splitting, E) modified watershed based splitting and F) different nuclear phenotypes (from different cell types).

we propose an iterative approach. We restrict the splitting process only to clustered cells and iteratively increase the allowed splitting length. First clustered nuclei are detected based on the area of the segments. All segments below a certain threshold are considered to be isolated nuclei, all other nuclei are processed by the constrained watershed splitting method. The allowed splitting length is set low, such that only nuclei which briefly touch are separated. Then small segments are added to the final segmentation result, while the clusters are processed again using the watershed algorithm, but with an increased threshold on the splitting length. This process is repeated until all segments have an area within the expected range. An example of this iterative procedure is shown in Fig. 6.12. The original image is shown in Fig. 6.12.A, while the initial segmentation result can be seen in Fig. 6.12.B. The grey segments correspond with segments whose area exceed a predefined threshold, typically corresponding with clustered nuclei. The remaining nuclei (shown in white) are considered isolated, and remain unchanged in the rest of the splitting algorithm. The grey segments are then split using the length constrained watershed separation algorithm. Fig. 6.12.C shows the splitting result when the maximal splitting length is set to 4 pixels. As a result of this splitting,

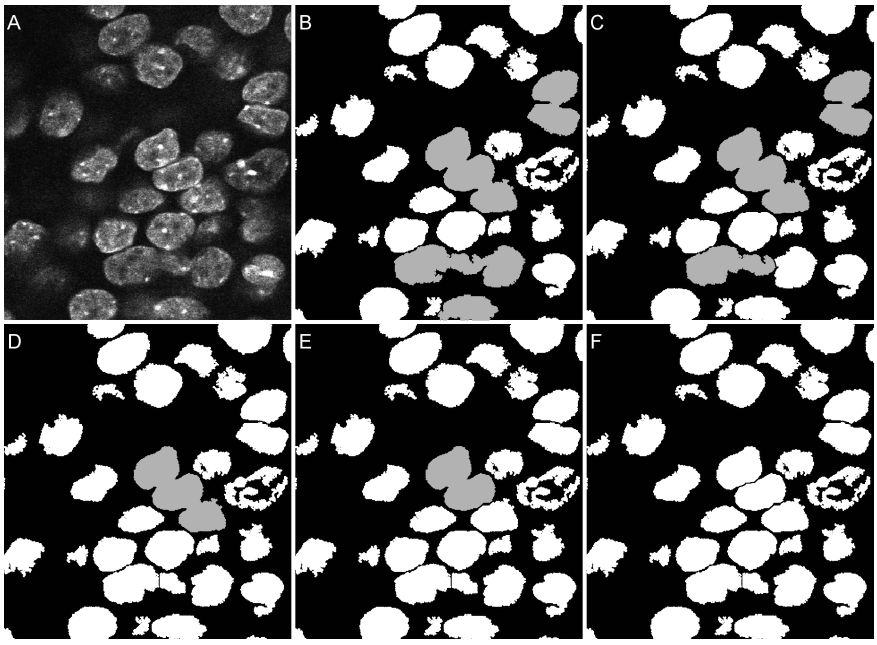


Figure 6.12: An illustration of the iterative splitting of nuclei. A) the original micrograph, B) the initial segmentation result, where the clustered cell are depicted in grey. C-F the result of the iterative splitting steps.

some nuclei are separated. These newly found segments are again checked to see if they exceed the area threshold. The remaining clustered nuclei are again processed by the splitting algorithm, but now with a splitting length of 6, 8 and 10 pixels. The results of these iterations are respectively shown in Fig. 6.12 D-F. By only processing likely clusters, the risk of over-segmentation is reduced, while large clusters are still separated, even if they touch for a longer distance, as can be seen in Fig. 6.12.F.

6.2.2.3 Error correction

Given the complexity in appearance and the density of nuclei in the brain, it should not be a surprise that not all nuclei are correctly detected, especially considering the fact that each slice is segmented independently. Therefore an extra step is proposed. This extra segmentation step attempts to detect and correct erroneous segments.

Segmentation quality control Previous studies in literature proved that it is possible to detect erroneous segments by analysing basic phenotyping features

[Chen 06, De Vos 10b, Ochoa Donoso 10, Ochoa Donoso 11]. Using these features, all segments are classified as correct or as potentially erroneous. The potentially erroneous segments will be corrected, using information from other slices. To distinguish between correct and erroneous segments, the following features are calculated for each segment:

- Area, the number of pixels in the nucleus.
- Perimeter, the distance around the boundary of the found nucleus.
- Solidity, the ratio between the area of the nucleus and the area enclosed by the convex hull of the nucleus [Jansen 09]. The convex hull of a segment corresponds to the contour with the shortest perimeter possible which envelopes all pixels from the nucleus.
- Major axis length, the length (in pixels) of the major axis of the ellipse that has the same normalized second central moments as the nucleus. See Chapter 8 for more detailed information.
- Minor axis length, the length (in pixels) of the minor axis of the ellipse that has the same normalized second central moments as the nucleus.
- The average intensity of the nucleus
- Clustered nuclei, a binary value expressing if there are other segments almost touching this segment, i.e. within a distance of 3 pixels.
- The number of segments found in the neighbouring slices at the same location as the current segment.

Based on the feature values of a set of manual classified segments, a random tree classifier was trained [Breiman 01, Hall 09]. Nuclei classified as erroneous segments are discarded, trusting that the next phase will detect a correct segmentation, based on prior knowledge coming from surrounding nuclei.

Handling missing slices So far the proposed segmentation algorithm only used 2D information, however for the difficult cases it might be beneficial to incorporate the segmentation result of surrounding slices. For each nucleus found in the previous slice, we check if a corresponding 2D segment is found in the current slice. Two segments are considered to correspond to the same nucleus if the overlapping area is above 50% of the area of the union of both segments. If no corresponding segment is found, the segment contour of the previous slice will be used as an initialization for finding the segment in the current slice. Updating an initial contour to delineate a segment is a well studied approach in literature, e.g.

using snakes [Kass 88, Xu 98], or active shape models [Cootes 95, Cootes 01], or dynamic programming like approaches [Sun 09, Cattrysse 13], etc.; however we use a simple, non regularized approach. This approach has the benefit of computational efficiency, while still achieving good results.

The segmentation of missing nuclei is inspired by the contour optimization discussed in [Cootes 95, Isard 98a]. We will illustrate the working of this step using Fig. 6.13. The proposed segmentation result is shown in Fig. 6.13.A. The detected nuclei are delineated in magenta. Two nuclei are not detected, as can be seen. At the location of these nuclei, a single segment was found in the previous slice. The contour of this segment is now superimposed on the current slice, shown in cyan. This contour is Gaussian downsampled such that only one out of twelve contour points remain. The goal is now to move this set of contour points, such that they coincide with the nuclear membrane, i.e. the physical boundary of the nucleus. In contrast to many active contour implementations, we restrict the possible displacements along the direction normal to the boundary. These boundary normals are illustrated in green. Along these normals, a crisp transition from dark to bright should correspond to the nuclear boundary. This corresponds to a negative derivative of the intensities along the normals, shown for some of the normals in Fig. 6.13.B. Note that for the derivative profile corresponding to the normal IV, that a part of the profile is set to 0. This is because that part already belongs to a different segment, and since a pixel can only belong to one nucleus, we restrict the possible locations for the contour to unoccupied pixels. The locations corresponding to the minima of the derivatives are back projected to the image, Fig. 6.13.C. In order to have a smooth nuclear contour, the contour is upsampled using cubic spline interpolation.

Note that there is still one nucleus undetected in the example of Fig. 6.13.C. This is because this was the first slice that the nucleus occurred, thus it was not present in the previous frame and as such was not detected. Since the slice-by-slice segmentation has most troubles with these "outer" slices, but can segment the nucleus robustly at the centre of the nucleus, this problem can be circumvented by first processing the 3D data stack starting from the top slice, propagating to the bottom slice and then repeating this process but in the opposite direction.

This approach finds the nuclear contour in the current slice, under the assumption that the boundary of the nucleus is in the close vicinity of the location of the nuclear boundary in the previous slice. Even if this assumption is violated, the method will typically still find a large part of the nucleus in the slice. This segment will generally be sufficient to correctly merge 2D segments into a 3D segment (see the following subsection). If desired, it is possible to improve the 2D segmentation acquired in this step using the shape of the full 3D segment. This would result in an extra postprocessing step. This extra step of course comes at an extra compu-

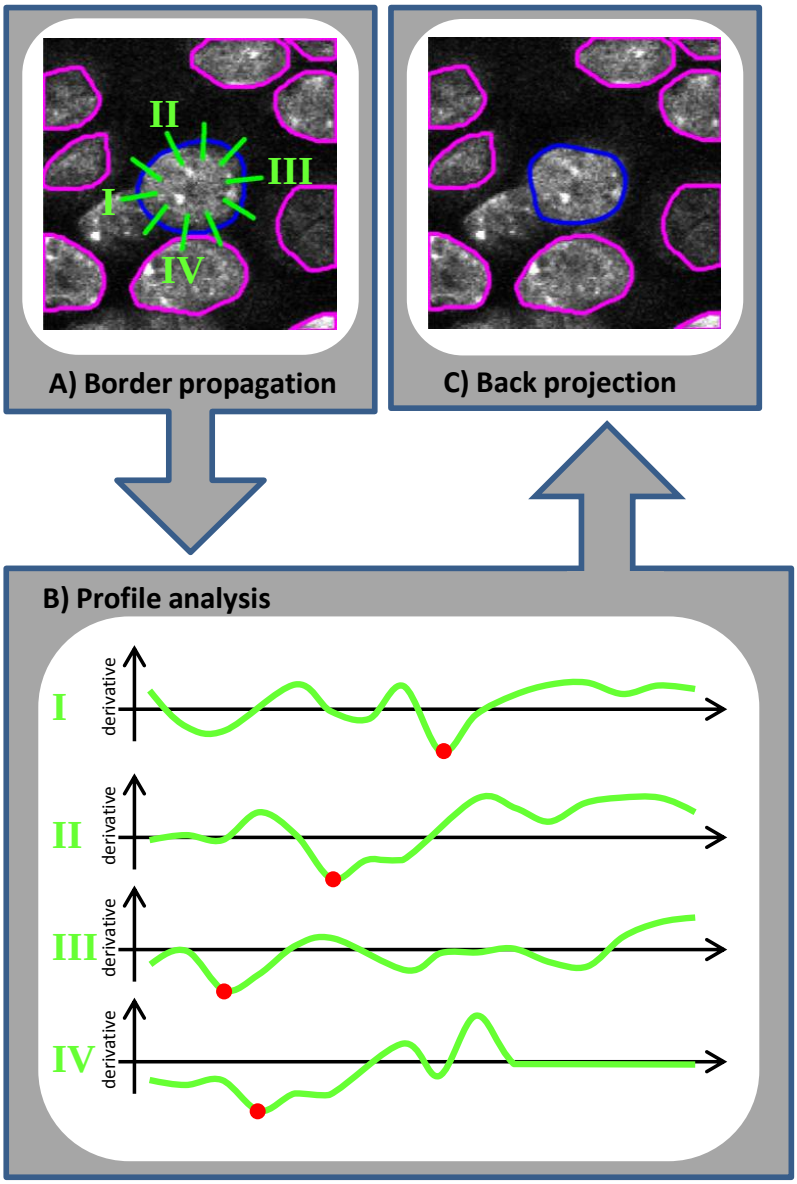


Figure 6.13: An overview of the contour propagation process. A) in blue the contour from the previous slice is shown, normal to this contour the profiles will be calculated (symbolized in green), in pink the found segments in the current slice are shown. B) the gradient along the profiles are shown, where the minima correspond to the nuclei contours, (shown as red dots). C) The found minima are back projected to there initial location. Based on these points a smooth continuous contour is calculated using b-spline interpolation (shown in blue).

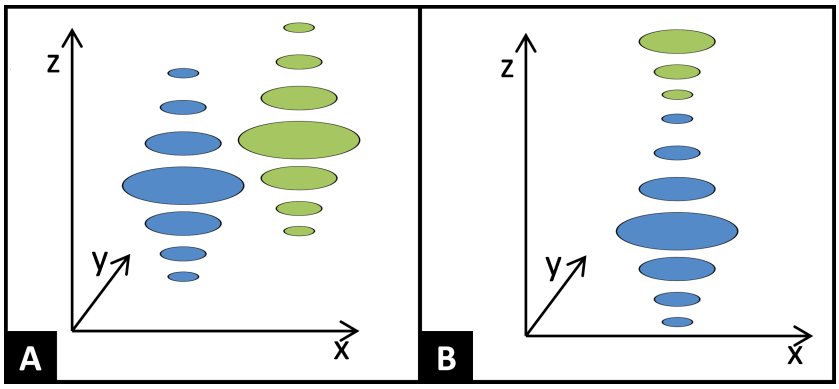


Figure 6.14: Two examples of nuclei configurations which have to be handled by the 3D merging step. A) two nuclei close to each other (coloured in green and blue) spread over 8 different slice, B) two nuclei with approximately the same centroid in the (x,y) plane and touching in the Z direction.

tational cost. For our purpose the gain is not worth the cost, therefore we won't elaborate on this strategy. For the interested reader we refer to a number of different strategies from literature that could be used for this purpose [Malladi 95, Barequet 96, Sethian 99, Jeong 07, Charmi 08].

6.2.2.4 3D Merging and Splitting

The previous steps result in a set of 2D segments. Combining these into 3D segments is a common problem in 3D segmentation for which rather straightforward solutions exist. The most widely used approach is based on the connected components algorithm [Vincent 91, De Bock 10], which assigns a label to each foreground pixel such that all touching foreground pixels have the same label assigned. A 3D segment then corresponds to all pixels with the same label. This however means that a 2D segment touching a segment in the consecutive slice will be merged, even if these two segments only touch at a single pixel location. This would result in the merging of almost all nuclei in the dentate gyrus, since this is a location in the brain where the nuclei are extremely densely packed, so almost all touch each other in 3D space. Hence we will apply a more restrictive approach, instead of merging all touching 2D segments, we propose to only merge segments which have an overlap of over 60% between consecutive slices, i.e. the area from the intersection of both segments corresponds to at least 60% of the area of the union of both segments. Figure 6.14.A illustrates this process. Segments of consecutive slices, shown as ellipses, are given the same label (blue or green) if and only if they have sufficient overlap.

It can occur that two 2D segments have sufficient overlap, while corresponding to different nuclei. An example of this is shown in 6.14.B, where a new nucleus is detected at the same location where in the previous slice a different nucleus was present. These situations are easily detected be exploiting the ellipsoid like shape of nuclei: the 2D area of a single nucleus will be highest in the slice through the centre of the nucleus and will decrease quasi monotonically towards the outer slices. So by detecting the slices where the area of a segment is smaller than the area in the surrounding slices, we find the location to split the 3D segment.

6.2.2.5 Experimental results

Due to the lack of accurate ground truth for 3D nuclear segmentation of fluorescence microscopic images of the brain, validation of our proposed method is not straightforward. To overcome the lack of real ground truth, we will validate our method using a synthetic data set. In this synthetic data set nuclei are approximated by an ellipsoid. The length of the semi-primal axes of all ellipsoids, $\mathbf{a}$, were modelled using a multivariate normal distribution: $\mathbf{a} \sim \mathcal{N}(\boldsymbol{\mu}, \Sigma)$, with $\boldsymbol{\mu} = (70, 50, 50)$ and $\Sigma = \mathrm{diag}(12, 7, 7)$. All nuclei are randomly spaced (based on a uniform distribution) and randomly oriented (also based on a uniform distribution) within a 3D voxel space, with the only constraints that the ellipsoids can not overlap. The intensity within an ellipsoid, i, is set as a constant value for each ellipsoid, with $i \sim \mathcal{N}(95, 14)$.

While the previous paragraph explains how the 3D nuclear configuration is simulated, this does not result in realistic looking images. Microscopy is hampered by a number of image degradations such as blur and noise. In order simulate realistic images, we model the degradations on our input image using the following equation:

$$f'(\mathbf{x}) = s((f * h)(\mathbf{x}) + b) + w(\mathbf{x}), \tag{6.2.7}$$

where b represents a constant background signal, for this simulation it was set at 7.8. The convolution, $(f * h)(\mathbf{x})$, is responsible for the blurring of $f(.)$, where $h(.)$ corresponds to a Gaussian kernel with $\Sigma = \mathrm{diag}(5, 5, 5)$. A realistic microscopic image is corrupted by two types of noise: shot noise and detector noise [Rooms 05b, Vonesch 13]. The shot noise is represented by $s(\lambda)$ which corresponds to a Poisson random variable of intensity λ, this noise is due to the quantum nature of light, while the detector noise is modelled using $w(\mathbf{x}) \sim \mathcal{N}(0, 3.1)$. All parameters of the degradation model are in accordance with the ISBI 2013 deconvolution challenge [Vonesch 13], which trained these parameters to simulate real microscopy image degradations. An example of a such a microscopic slice is shown in Fig. 6.15.

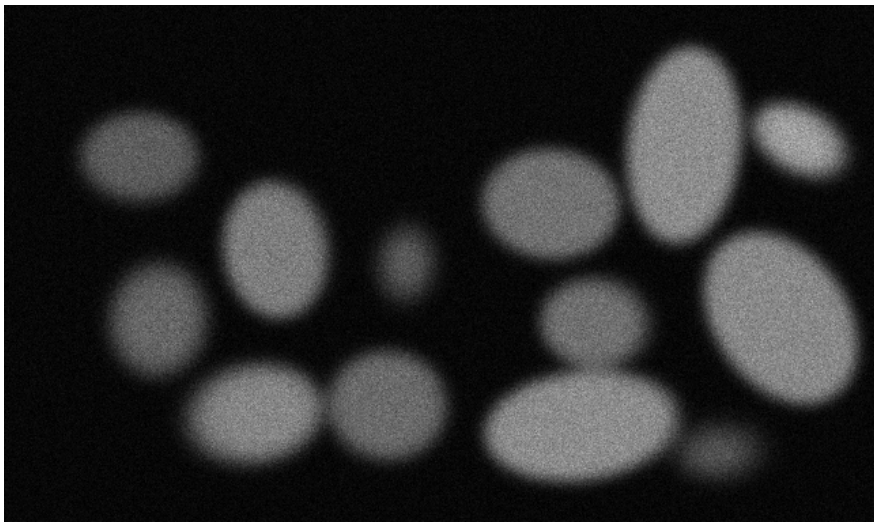


Figure 6.15: A synthetic microscopic slice showing clustered nuclei.

While an abundance of open access software exists for the analysis of 2D microscopic images, fully 3D segmentation tools are rather rare. We compared our method with the results from CellSegm [Hodneland 13], the level set based approach from Chinta [Chinta 12] and the commercial software package Huygens [Huygens 11]. The segmentation results for this synthetic data set are summarized in Table 6.4 (for an explanation on the used metrics we refer to appendix B). CellSegm assumes that the majority of the nuclei lie in the same plane. This assumption is exploited in the splitting step, where the segmentation result of a single predefined slice is used to split touching nuclei in all other slices. While this is a reasonable constraint for many in vitro microscopic applications, it is too restrictive for in situ imaging where nuclei are organized in a seemingly random configuration. This results in erroneously splitting of several nuclei, which results in a high number of false positives and in a decrease in Dice coefficient. The method proposed by Chinta et al. is more robust against over-segmentation. Although the method has some problems detecting nuclei with a lower average intensity, resulting in two false negatives. The method is also slightly hampered by blur: due to blur, a light grey hallo surrounds bright nuclei, this hallo is occasionally also detected as a segment. These false positives might be removed using an extra classification step based on shape parameters. The noise and blur also result in slight errors on the location of segment boundaries, explaining the slightly lower Dice coefficient.

The Huygens software segments solely based on intensity, without any constraint on shape or splitting strategy. This results in a trade-off between under segment-

method	TP	FN	FP	P	R	F	Dice
Hodneland et al.	13	4	15	0.464	0.765	0.578	0.791
Chinta et al.	15	2	2	0.882	0.882	0.882	0.891
Huygens software	11	6	1	0.917	0.647	0.759	0.734
proposed	**17**	**0**	**0**	**1**	**1**	**1**	**0.957**

Table 6.4: Comparison of the detection and segmentation results from different state-of-the-art methods for 3D nuclei in a synthetic data set (TP = true positives, FN = false negatives, FP = false positives, P = precision, R = recall and F = F-measure).

ation and completely missing nuclei. The software gives adequate results for slightly touching nuclei, but is not robust against densely clustered nuclei. For population statistics this does not have to be an issue, since irregular segment phenotypes can be discarded using a number of available post processing filters. However for applications where the configuration of individual nuclei is paramount, other methods are more appropriate. Our proposed method is able to detect all individual nuclei, without finding erroneous detections. Not only the detections but also the actual segmentation is accurate, i.e. a Dice coefficient of 0.96.

Fig. 6.16 shows an example of the segmentation result of a real microscopic image. This micrograph shows a cross section of the dentate gyrus, a part of the hyppocampus where nuclei are densely clustered, i.e. about 600 nuclei in a volume of $128 \times 212 \times 40$ micron³. Fig. 6.17 shows a 3D rendering of our segmentation output. With this example image both the nuclei density and the complexity of the spatial configuration of cell nuclei becomes apparent. As can be seen in Fig. 6.16, all nuclei, irrespective of the variability in phenotypes, are correctly detected. Even nuclei with low average intensity or with elongated small shapes where correctly found without over-segmenting them.

In Fig. 6.18 a volume histogram of the segments found by the different methods is shown. As a reference ground truth histogram a set of isolated nuclei where selected from the top of the neocortex and neuroglia. Note that this ground truth consists of multiple cell types and hence contains nuclei with varying phenotypes. This explains the multi-modal histogram. This reference histogram was calculated based on 342 segments. Although the cell types in the validation dataset consist mainly of a single phenotype, this ground truth still contains some valuable information: 1) few nuclei have a value below 5000 voxels (the few that do, is because they are at the border of the image and their total volume can not be estimated), 2) The large majority of cell nuclei have a volume less than 18000 voxels. Experts expect that the volume of cell nuclei of the same cell type are uni-modally distributed.

Both cellSegm and Huygens have problems of over segmentation, which can be seen by the peaks at volumes under 5000 voxels. Huygens also had a problem

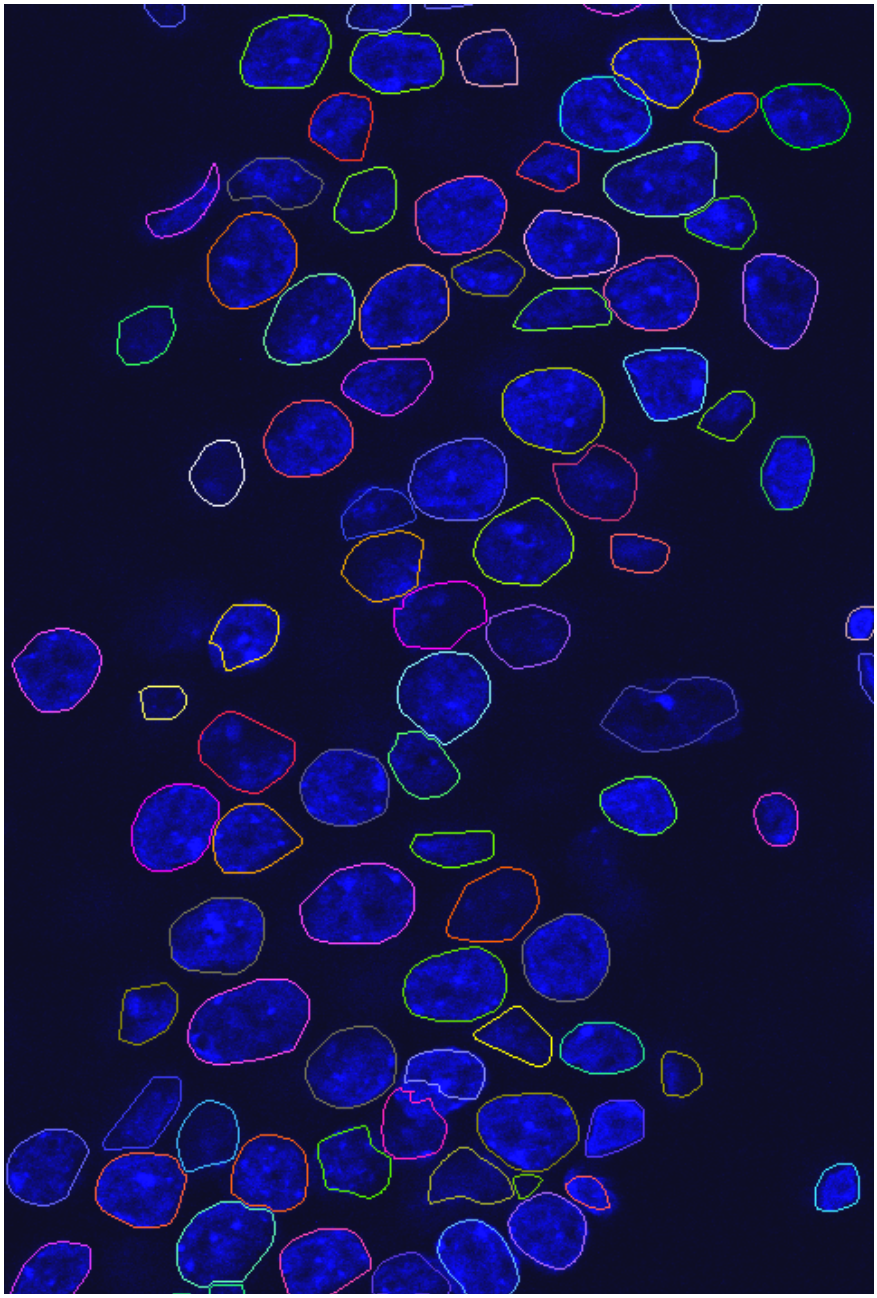


Figure 6.16: An example of a slice segmented using the proposed method. In blue the original DAPI stained micrograph is shown. The segments are depicted by their contour which are superimposed using a different colour for each segment.

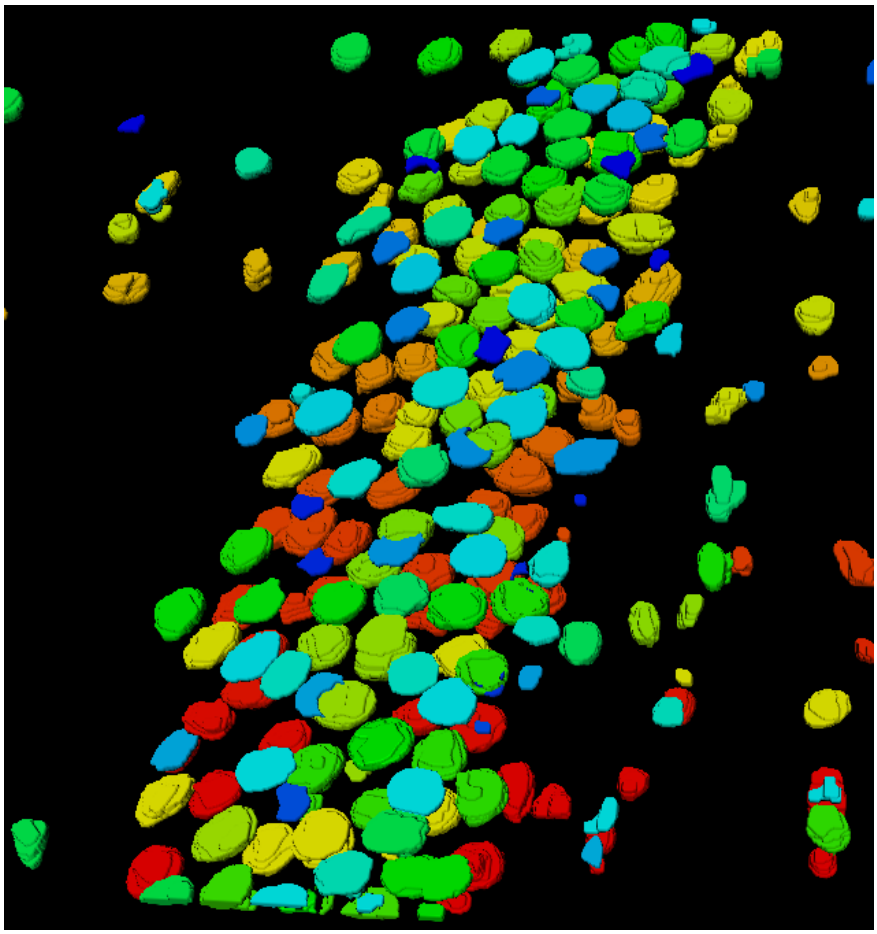


Figure 6.17: An 3D rendering of the segmentation output. The dense band of nuclei corresponds to dentate gyrus, a part of the hyppocampus.

of under segmentation, it found a single segment of over 400000 voxels, which is not shown in the histogram. This single segment corresponds to approximately 360 nuclei, based on the average nuclear volume of the reference histogram. The dentate gyrus, i.e. the brain section analysed in this histogram, is mainly dominated by a single cell type/phenotype, this corresponds to the approximately uni-modal distribution resulting from the segmentation results generated by the proposed method. The proposed method found 629 nuclei in this data set, while both Huygens and cellSegm both found only less than 90 nuclei.

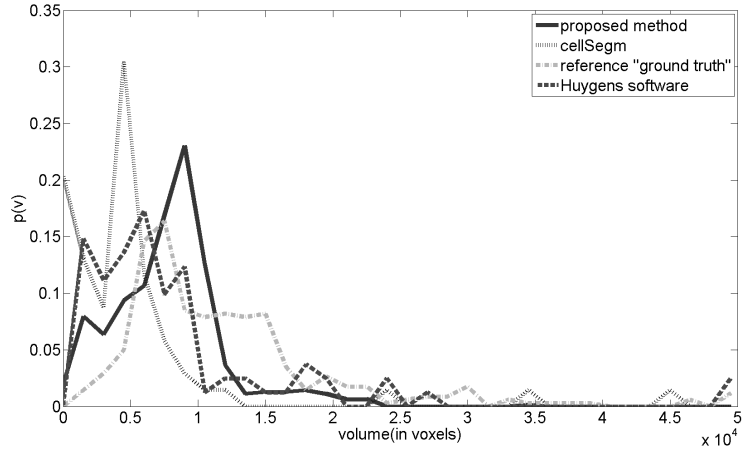


Figure 6.18: A volume histogram of the found segments in the dentate gyrus of the brain.

6.2.3 Use case: peripheral blood mononuclear cell nuclei

While the previously proposed method solves the segmentation of clustered nuclei without putting hard constraints on the type of nuclear configuration, there are still applications for which the proposed algorithms is of little use. This is mainly because it requires 3D information, whereas for lot of pharmaceutical applications, 2D microscopic images are used. Such a pharmaceutical application is the study of autoimmune disease. Autoimmune diseases affect approximately 5–8% of the human population and of all subjects with an autoimmune disease, 78% are women [Jacobson 97]. An explanation might be found in the presence of fetal cells in the maternal circulation and tissues [Lepez 11]. During pregnancy, fetal cells cross the placenta into the maternal circulation. Fetal cells can persist in maternal blood, developing into micorchimersims. The persistence of these fetal cells may result in the development of autoimmune diseases, such as autoimmune thyroid diseases (AITD). In the case of fetal microchimerism in autoimmune diseases,

automatic detection of fetal peripheral blood mononuclear cells (PBMCs) in the maternal circulation can shed light on the potential role of fetal cells in these diseases [Lepez 11]. It is however very labour-intensive to locate one or more fetal cells in a population of millions of maternal cells. Staining of the male fetal cells using Fluorescence in Situ Hybridization (FISH) with subsequent DAPI staining and automatic detection of all cell nuclei on a slide, can facilitate the detection of fetal cells.

Most nuclei segmentation approaches, both edged based and intensity based, are non optimal in the sense that the detection and recognition steps are treated independently: the splitting into individual nuclei only uses the result of the initial segmentation, instead of all available information, i.e. the complete image. Thus, incorrectly segmented pixels or edges can have a big influence on the final detection of individual nuclei. Because both tasks can be mutually beneficial [Horvath 09, Nemeth 11], we propose a novel method which jointly optimizes the foreground, i.e. pixels of interest, classification and the separation of segments into individual nuclei. The decision of segmenting a pixel as foreground is based both on intensity and on the likelihood that the pixel is part of a nucleus. Such a joint optimization has been proposed in [Horvath 09, Nemeth 11], where individual circular objects are modelled using Markov random fields. The approach shows good results for circular objects of the same size, but analysing objects with different sizes is non-trivial. Furthermore, their joint model is optimized using meta-heuristics such as simulated annealing, so it is dependent of the initialization of the optimization process. In this work we propose a different model which also jointly optimizes the detection of relevant pixels and the detection of individual nuclei. The proposed method fits within the general active contour framework (see chapter 5), but with a novel shape prior specifically developed for the detection of cell nuclei. We propose a shape regularization term that exploits the regular shape of cell nuclei, penalizing segments which strongly deviate from the expected shape. In this work, we model a nucleus as a disk. This allows us to build a dictionary consisting of binary images that correspond to the segmentation result of a single nucleus with a predefined radius and location, called atoms. Any segmentation result of an image containing multiple nuclei, can be approximated linearly as a superposition of atoms. We will use this approximation as a new regularization term in the active contour framework. The segmentation result is calculated by minimizing a convex energy function, such that the active contour is invariant to initialization [Chan 06, Bresson 07a]. The proposed joint convex optimization approach results in accurate nuclei detection, especially for realistic segmentation problems, where cell nuclei can have different sizes.

6.2.3.1 Joint shape and intensity constraints

Let S be the set of all functions of the form $\Omega \mapsto [0,1]$, where $\Omega = [1,m] \times [1,n]$, and $m \times n$ the dimension of the image. We will represent the segmentation result as a characteristic function $u(.) \in S$. A pixel $\mathbf{x}$ that corresponds to a nucleus, will be represented by one, whereas background pixels are represented by zero, i.e.

$$u(\mathbf{x}) = \begin{cases} 1 & \text{if } \mathbf{x} \text{ represents a nucleus pixel,} \\ 0 & \text{otherwise.} \end{cases} \tag{6.2.8}$$

In chapter 5, the segmentation process was described as the calculation of a MAP estimator. This MAP estimator corresponds to the following optimization problem:

$$\hat{u}(.) = \underset{u(.) \in S}{\arg\min} \left| \nabla u(.) \right| + \beta \langle u(.), r(.) \rangle, \tag{6.2.9}$$

with

$$r(\mathbf{x}) = (\mu_f - f(\mathbf{x}))^2 - (\mu_b - f(\mathbf{x}))^2, \tag{6.2.10}$$

where $f(.)$ corresponds to the image, μ_f, μ_b are respectively the expected foreground and background intensity and β is a weighting parameter used to tune the influence of the data-fit term in relation to the total variation regularization. In this MAP estimator total variation is used to remove noisy segmentation pixels. This regularization is useful if pixels are incorrectly classified, i.e. background pixels detected as foreground or vice versa, due to noise in the image. In microscopic images however, incorrectly detected nuclei are often caused by clutter in the image, e.g. dead cells or bleed-through from other fluorescent channels. This is not solved using generic priors since these incorrectly detected nuclei are natural objects, e.g. they have smooth boundaries. Therefore, a number of shape based regularization terms have been proposed [Schoenemann 07, Charmi 08, Cremers 08, Klodt 11, Liu 12]. However these shape priors are limited to images with only a single object of interest with a specific shape. We propose a regularization term that exploits the regular shape of cell nuclei, penalizing segments which strongly deviate from the expected shape, while not constraining the number of nuclei.

In this work, a nucleus is modelled as a disk. For a given radius, r, and location, $\mathbf{y}$, we can calculate the ideal $u(.)$, i.e. a binary image where the pixels within a distance r of $\mathbf{y}$ are equal to one and all other pixels are equal to zero:

$$d(\mathbf{x}; \mathbf{y}, r) = \begin{cases} 1 & \text{if } \|\mathbf{x} - \mathbf{y}\|_2 < r \\ 0 & \text{otherwise} \end{cases}$$

This is of course under the assumption that there is only one nucleus in the image. We will refer to each of these possible segmentation results as *atoms*. In most applications however the image does contain multiple nuclei. Even the number of nuclei is typically not known. Therefore we model the unknown segmentation $u(.)$ as a superposition of disks, i.e.

$$u(.) \approx \sum_{\mathbf{y}} t(\mathbf{y}) d(.; \mathbf{y}, r) \tag{6.2.11}$$

where $t(\mathbf{y})$ represents a characteristic function expressing the presence of a nucleus centred at location $\mathbf{y}$. Note that this representation expects a predefined diameter of a cell nucleus. However the size of a nucleus is generally not fixed, but can be considered to lie in an interval $[a, b]$. Therefore we can extend eq. (6.2.11) to approximate $u(.)$ using a dictionary of atoms corresponding to discrete range of nuclei sizes:

$$u(.) \approx \sum_{r=a}^{b} \sum_{y} t(\mathbf{y}; r) d(.; \mathbf{y}, r) \tag{6.2.12}$$

with $t(\mathbf{y}; r)$ corresponding to the presence of a nucleus of size r at location $\mathbf{y}$. Note that this linear combination penalizes overlapping nuclei, which is desirable since a pixel can only correspond to a single nucleus. A good segmentation is one which consists of a small number of atoms. This sparsity constraint can be used as a new regularization term:

$$\big(\hat{u}(.), \hat{t}(.;.)\big) = \underset{\substack{u(.) \in S \\ t(.;.) \in S^e}}{\arg\min} \sum_{r} \Big\| t(., r) \Big\|_0 + \beta \big\langle u(.), c(.) \big\rangle \tag{6.2.13}$$

such that

$$u(.) \approx \sum_{r} \sum_{\mathbf{y}} t(\mathbf{y}; r) d(.; \mathbf{y}, r) \tag{6.2.14}$$

with e the number of disc sizes considered, i.e. $e = b - a$. The energy term minimized by $u(.)$ is based on a l_0 norm which is defined as $\big\|u(.)\big\|_0 = \sum_{\mathbf{x} \in \Omega} 1 - \delta_{0, u(\mathbf{x})}$, where $\delta_{i,j}$ represents the Kronecker delta, which is equal to one if i and j are equal and is equal to zero in all other cases. This l_0 norm regularization comes down to calculating the number of nuclei, i.e. the number of non-zero elements in $t(.;.)$. The l_0 norm is non-convex, hampering optimization. Fortunately the l_0 norm can be approximated by the l_1 norm which is the closest convex norm to l_0. It has been proven that this approximation of an l_0 norm gives good results for many different applications, e.g. for compressed sensing, sparsity based image restoration, dictionary based super resolution, etc. [Baraniuk 07, Baraniuk 10]. This new prior

results in the following active contour:

$$\big(\hat{u}(.),\hat{t}(.;.)\big) = \underset{\substack{u(.)\in S\\ t(.;.)\in S^e}}{\arg\min} \sum_r \Big|t(.,r)\Big| + \beta\big\langle u(.),c(.)\big\rangle \tag{6.2.15}$$

such that eq. (6.2.14) is met. In order to optimize the constrained problem in eq. (6.2.15) the problem is approximated by adding the constraint in the form of a quadratic term, resulting in the following unconstrained optimization problem:

$$\begin{aligned}\big(\hat{u}(.),\hat{t}(.;.)\big) \;=\; & \underset{\substack{u(.)\in S\\ t(.;.)\in S^e}}{\arg\min} \sum_r \Big|t(.,r)\Big| + \beta\langle u(.),c(.)\rangle \\ & +\gamma\Big\|u(.) - \sum_r\sum_{\mathbf{y}} t(\mathbf{y};r)d(.;\mathbf{y},r)\Big\|_2^2,\end{aligned} \tag{6.2.16}$$

where γ is a weighting parameter. Note that this only approximates the constraint in eq. (6.2.12). Although there exists efficient techniques to enforce this constraint exactly, e.g. augmented Lagrangian or ADMM methods, we propose to use the approximation in eq. (6.2.16) instead. This allows the active contour to detect nuclei whose shape slightly deviates from the circular model or to detect partially overlapping nuclei. Given the convexity of eq. (6.2.16), this problem can be solved by iteratively optimizing for $\mathbf{t}$ and $\mathbf{u}$ independently, i.e.

$$t(.;.) \;=\; \underset{t(.;.)\in S^e}{\arg\min} \sum_r \Big|t(.;r)\Big| + \gamma\Big\|u(.) - \sum_r\sum_{\mathbf{y}} t(\mathbf{y};r)d(.;\mathbf{y},r)\Big\|_2^2, \tag{6.2.17}$$

$$\hat{u}(.) \;=\; \underset{u(.)\in S}{\arg\min}\, \beta\langle u(.),c(.)\rangle + \gamma\Big\|u(.) - \sum_r\sum_{\mathbf{y}} t(\mathbf{y};r)d(.;\mathbf{y},r)\Big\|_2^2. \tag{6.2.18}$$

The problem in eq. (6.2.17) can be solved using Newton's method, which iteratively updates $t(\mathbf{x};r)$ using the following scheme:

$$\begin{aligned}t^{(2k+1)}(\mathbf{x};r) \;=\; & t^{(2k)}(\mathbf{x};r) \\ & -\frac{\frac{1}{\gamma} - 2\big(u(\mathbf{x}) - \sum_r\sum_{\mathbf{y}} t^{(2k)}(\mathbf{y};r)d(\mathbf{x};\mathbf{y},r)\big)\sum_{\mathbf{z}} d(\mathbf{z};\mathbf{x},r)}{\big\langle d(.;\mathbf{y},r), d(\mathbf{x};.,r)\big\rangle},\end{aligned} \tag{6.2.19}$$

where the subscript index represents the iteration step. The constraint that $t(\mathbf{x};r) \in [0,1]$ can be softly enforced by adding a barrier function to eq. (6.2.17). For this purpose we propose the use of a piecewise linear barrier function (discussed in more detail in chapter 5):

$$b(x) = \begin{cases} -1000\,x & \text{if } x < 0, \\ 1000\,x - 1000 & \text{if } x > 1, \\ 0 & \text{otherwise.} \end{cases} \tag{6.2.20}$$

This barrier function has the benefit of penalizing values outside of the interval $[0, 1]$, while not giving preference to any specific value inside the interval. The piecewise linear barrier function is prone to overshooting using the Newton-Raphson method, i.e. if $\mathbf{t}_k$ is greater than 1, then $\mathbf{t}_{k+1}$ would be equal to $-\infty$, whereas if $\mathbf{t}_k$ is less than 0, it would result in $+\infty$. However given the specific nature of this barrier function, with the minima corresponding to the roots of the function, it is possible to minimize this by Newton's for finding the roots of a function. This results in the following update step:

$$t^{(2k+2)}(\mathbf{x}; r) = t^{(2k+1)}(\mathbf{x}; r) - b'(t^{(2k+1)}(\mathbf{x}; r)) \tag{6.2.21}$$

with

$$b'(i) = \begin{cases} -i & \text{if } i < 0, \\ i - 1 & \text{if } i > 1, \\ 0 & \text{otherwise.} \end{cases} \tag{6.2.22}$$

The optimal $u(.)$ in eq. (6.2.18) can be found by solving a set of Euler-Lagrange equations. For an optimal $u(.)$, the following optimality condition should be satisfied:

$$u(.) = -2\frac{\beta}{\gamma}c(.) + \sum_r \sum_{\mathbf{y}} t(\mathbf{y}; r) d(.; \mathbf{y}, r). \tag{6.2.23}$$

The solution of eq. (6.2.23) is unconstrained, i.e. $u(\mathbf{x})$ does not have to lie in the interval $[0, 1]$. However minimizing eq. (6.2.18) in $u(\mathbf{x})$ is equivalent to minimizing a quadratic function. So if $u(\mathbf{x}) \notin [0, 1]$ then the constrained optimum is either 0 or 1, since a quadratic function is monotonic in an interval which does not contain its extremum. So the constrained optimum is given by:

$$\hat{u}(\mathbf{x}) = \max\Big(\min\big(u(\mathbf{x}), 1\big), 0\Big). \tag{6.2.24}$$

6.2.3.2 Nuclei detection

By adding a sparsity constraint on $t(.;.)$, we penalize representations which use more atoms than necessary. The function $t(.;.)$ can be used to detect the centroids of the cell nuclei. However by using the l_1 norm as an approximation of the l_0 norm it is not possible to get the nuclei centroids explicitly from $t(.;.)$. This is because the l_1 norm does not favour binary functions, whereas the l_0 norm does. For

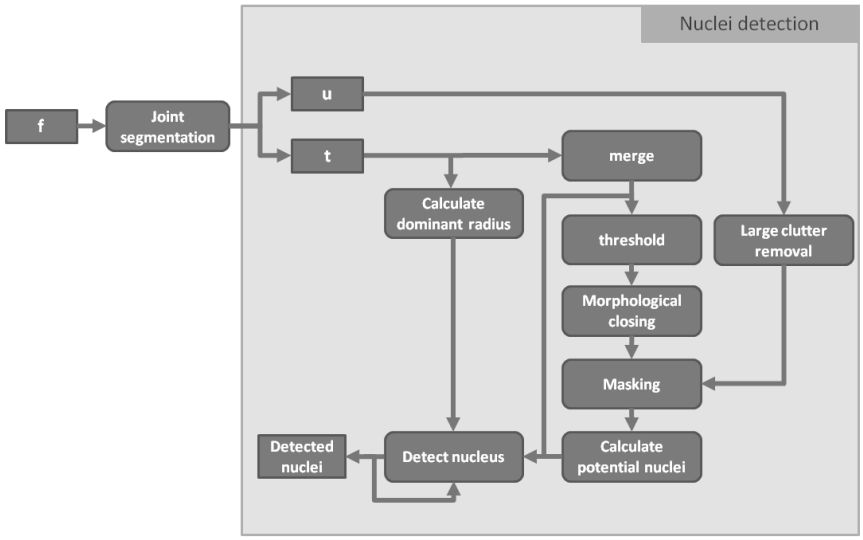


Figure 6.19: The work flow of the detection algorithm [De Vylder 13a].

example, consider a function $f_1(.)$ that has a zero value for its complete domain, except for a single point, where it has a value of one. If we compare this with a constant function $f_2(.)$, which for each argument returns a constant number such that $\int f_2(x)\,dx = 1$. While the l_1 norm of both functions is equal, this is clearly not the case for the l_0 norm, which is significantly lower for the binary function $f_1(.)$.

Due to the fact that t(.;.) is not binary, a special algorithm is required to detect the actual nuclei. A schematic overview of the proposed detection method is shown in Fig. 6.19. Since disks of multiple sizes can add to the representation of a nucleus in the segmentation result, we first detect the dominant radius for each pixel. Each pixel has a number of corresponding atoms, each representing a nucleus of different size. The radius corresponding with the atom with the highest value for a specific pixel is considered the dominant radius of that pixel, i.e. $\arg\max_r\ t(\mathbf{x}; r)$. Since the most important information is the location of a cell nucleus, a merging step is applied such that all $t(\mathbf{x}; .)$ weights corresponding to the same pixel location are combined by summing them together. This is done in parallel with the calculation of the dominant radius. The merging step results in a new image, $t_m(\mathbf{x}) = \sum_r t(\mathbf{x}; r)$, where high intensities occur at the centres of nuclei (shown in Fig. 6.20.D). In order to get crisp detections, the new image is converted to a

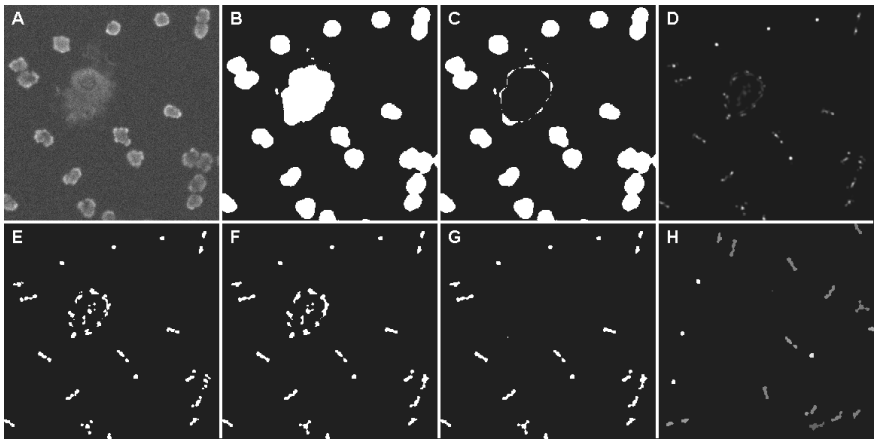


Figure 6.20: The different steps in the detection algorithm: (A) the original image, (B) the segmentation result, (C) large clutter removal, (D) the merged values for t, corresponding to the centroids of the disks used for the reconstruction of the segmentation result, (E) a binarization of (D), (F) show the morphological processing of (E), (G) is the masked version of (F) with (C) used as mask, (H) represents the intensity of each connected component by it's score, i.e. a measure of likeliness that that component corresponds to the centroid of a nucleus [De Vylder 13a].

binary image by thresholding it (Fig. 6.20.E). Due to small deviations in the shape of cell nuclei, it is possible that a single nucleus corresponds to multiple connected components in this binary image. Since these components are located in each other's vicinity, it is possible to overcome this problem by applying a morphological closing (Fig. 6.20.F). False detections due to noise and clutter are removed by merging the binary image with a mask, i.e. all pixels which are zero in the mask are set to zero (Fig. 6.20.G). The mask itself is calculated from the segmentation result $\hat{u}(.)$ (Fig. 6.20.B). First large clutter such as apoptotic cell nuclei, i.e. dead nuclei, are removed from $\hat{u}(.)$, by applying a morphological opening, and then using the inverse of this opened image as a mask on $\hat{u}(.)$ (Fig. 6.20.C).

Next, potential nuclei are detected by applying the connected component algorithm on the binary image. Finally, for each potential nucleus, i.e. connected component, a score is calculated. This score corresponds to the mean intensity in the preprocessed image of the pixels corresponding to the connected component in the $t_m(.)$ image (Fig. 6.20.H). The potential nucleus with the strongest score, $\mathbf{p}$, is removed from the potential list and is added to the list of detections. Since nuclei can only touch each other (they can not occupy the same space), all potential nuclei within a distance of the dominant radius of $\mathbf{p}$ can be removed from the list of potential nuclei. The steps of adding the strongest potential nucleus to the list of detections and removing "too close" potential nuclei are repeated until the list of potential

	RMSE	σ	**P**	**R**	**F**	**Dice**
EdgeProp	1.75	1.78	0.9973	0.9968	0.9971	0.941
CellProfiler	3.79	2.16	0.9914	0.9810	0.9862	0.939
CellC	16.24	3.63	0.9445	0.9945	0.9689	0.929
Proposed	**0.22**	**0.22**	**0.9998**	**1.0000**	**0.9999**	**0.981**

Table 6.5: A comparison of different cell nuclei detection and segmentation methods on the BBBC004v1 data set from the Broad Bioimage Benchmark Collection (RMSE= root mean squared error on the count of found cells per image, σ = standard deviation on the count of found cells per image, P = precision, R = recall, F = F-measure).

nuclei is empty. While the proposed detection algorithm lacks a strong theoretical foundation, it has the benefit of computational efficiency, while still achieving adequate results.

6.2.3.3 Experimental results

As a first validation of the proposed method, a synthetic data set is analysed [Ruusuvuori 08]. These synthetic images were proposed as a common benchmark for nuclei segmentation and detection. In Fig. 6.21(A) an example of such a synthetic raw image is shown. For these experiments, the parameters for eq. (6.2.13) are empirically chosen: $\beta = 5$, $\gamma = 1$, with a disk diameter equal to 20 pixels. The expected intensities needed in eq. (6.2.15) are estimated based on the result of a simple Otsu thresholding [Otsu 76, Sezgin 04]. Fig. 6.21(B) shows the detection result using the edge based detection method from section 6.1. The result of cellProfiler, which is a very popular example of an intensity based method [Kamentsky 11], is shown in Fig. 6.21(C), whereas Fig. 6.21(D) depicts nuclei detection using the proposed method. Correctly detected nuclei are shown with a green marker, errors are marked with an arrow: false detections with a red marker and cell nuclei which were not detected by the method are shown using a yellow marker. While edgeProp is able to separate clusters of nuclei, the method is limited by the strength of the edges, resulting in several nuclei which are not detected. CellProfiler, an intensity based method, is able to detect all nuclei clusters, but is not always able to separate them into individual nuclei. This can be seen by the occurrence of a red marker, a false detection, in the vicinity of multiple yellow markers, i.e. undetected nuclei. Clearly, the proposed method is more robust for detecting clusters of nuclei, while remaining able to detect all nuclei, even those with weak edges.

To quantitatively validate the result, a data set of 20 images, each containing 300 nuclei, is analysed. The results are shown in Table 6.5. The first row represents the results of an edge based method [De Vylder 11e], whereas the next two rows

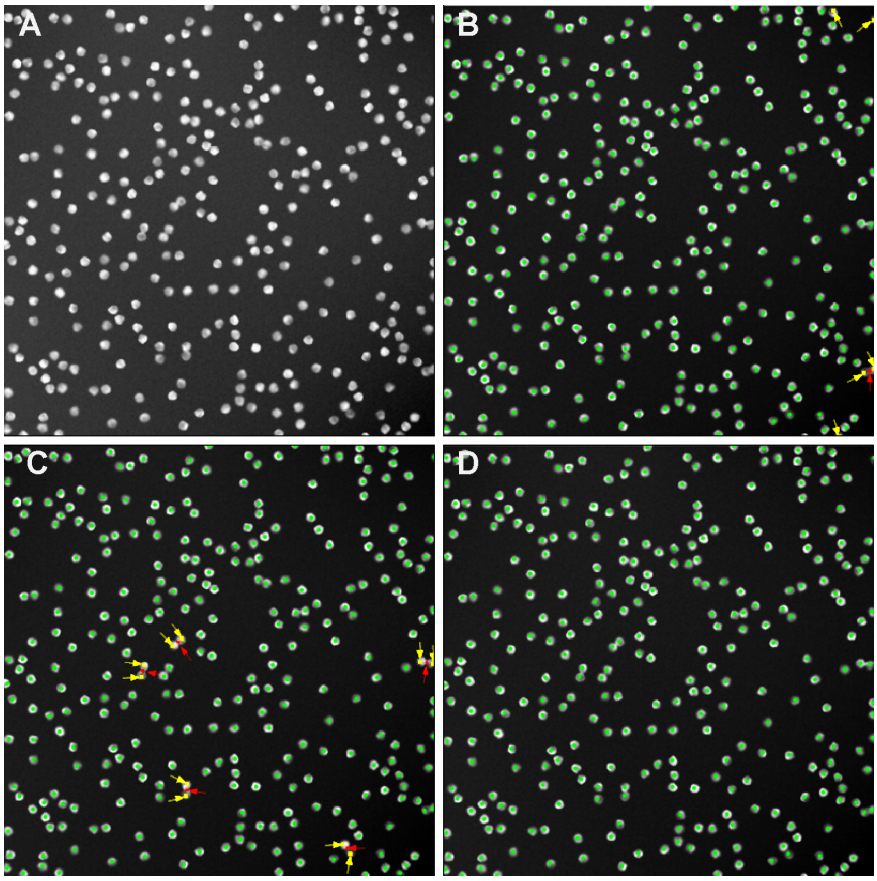


Figure 6.21: An example of cell nuclei detection using different methods on a benchmark image. (A) the original raw micrograph, (B)-(D) show the respective detection results superimposed on the preprocessed images. (B) edgeProp, (C) cellProfiler, (D) the proposed method. Correctly detected nuclei are shown with a green marker, errors are marked with an arrow: false detections with a red marker and cell nuclei which were not detected by the method are shown using a yellow marker. While edgeProp is able to separate clusters of nuclei, they are limited by the strength of the edges, resulting in several nuclei which are not detected [De Vylder 13a].

	data set 1	**data set 2**	**data set3**
# nuclei	3178	3197	582
# clusters	1025	1861	436

Table 6.6: ground truth statistics of the different data sets.

correspond with two intensity based images [Selinummi 05, Kamentsky 11]. The proposed method, presented in the last row, shows the best results for cell detection metrics (first 5 columns) as well as for the Dice coefficient. Note that lacking ground truth for individual nuclei, the Dice coefficient measures the similarity of all segmented nuclei compared with the ground truth for all nuclei instead, i.e. this is the Dice coefficient between two binary images, instead of between individual segments. While state-of-the-art methods perform reasonably for the nuclei count and detection part, the Dice coefficient shows still room for improvement. These good results are also due to the nature of this data set, which nicely models microscopic image degradations, but where the objects are not hampered by biological clutter, such as dead cells. A second validation is done using real microscopic images of nuclei from peripheral blood mononuclear cells (PBMCs) (see appendix A for more details). PBMC's consist of a number of different cell types, each with different nuclei morphologies: 75% lymphocytes (T cells, B cells, NK cells), monocytes, macrophages (rarely), basophils, dendritic cells, neutrophils (horse-shaped nucleus), eosinophils,...). Ground truth for these images is generated by manually annotating the data sets. This annotation was done by ourselves, and afterwards the annotations were checked and when necessary corrected by experts in the field of pharmaceutics. In Fig. 6.22 an example of nuclei detection in such a real microscopic image is shown. Fig. 6.22.A shows the ground truth detection, i.e. the manual annotations. Fig. 6.22.B and Fig. 6.22.C correspond to state-of-the-art detection methods [Kamentsky 11,De Vylder 11e]. Note that both methods erroneously detect nuclei at places where there is some smeared staining. When a nucleus ruptures it releases its staining, which results in bright smears. CellProfiler does not only suffer from false detections, but also merges touching nuclei. The proposed method is more robust against dye smears, while still being able to detect touching nuclei, as can be seen in Fig. 6.22.D.

Three different data sets were analysed, each data set with a different density of cell nuclei (Table 6.6). The first data set contained approximately 30% of touching nuclei, the second data set with almost 60% of the nuclei touching each other and finally a dense data set with 75% of clustered nuclei. For these experiments, the parameters for eq. (6.2.13) are empirically chosen: $\beta = 1$, $\gamma = \frac{1}{25}$, with a diameter in the range of [7.04μm-12.16μm]. The expected intensities needed in eq.(6.2.15) are estimated based on the result of a simple Otsu thresholding [Otsu 76, Sezgin 04]. State-of-the-art methods use the same prior knowledge, i.e.

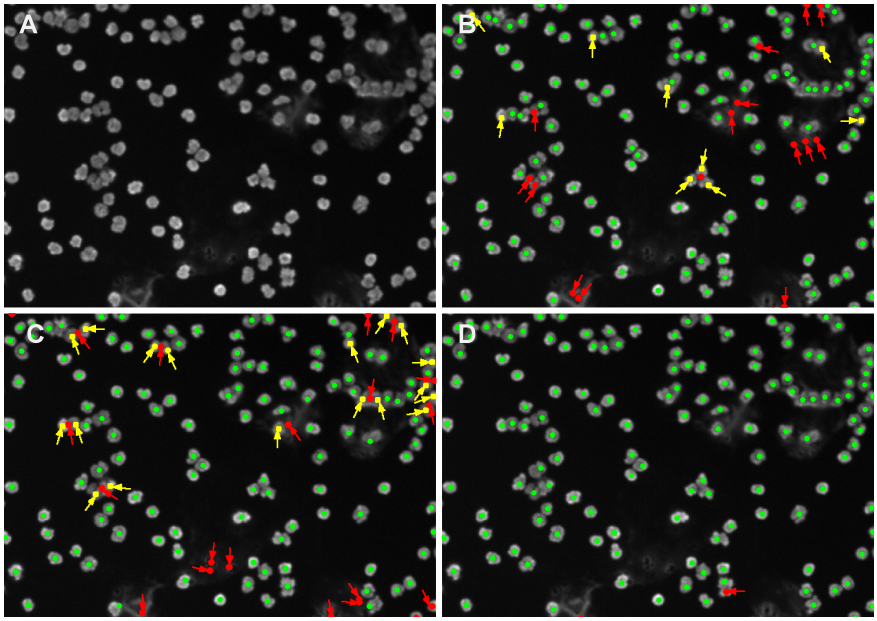


Figure 6.22: Example of cell nuclei detection using different methods on an image of peripheral blood mononuclear cells. (A) the original raw micrograph, (B)-(D) show the respective detection results superimposed on the preprocessed images. (B) edgeProp, (C) cellProfiler, (D) the proposed method. Correctly detected nuclei are shown with a green marker, errors are marked with an arrow: false detections with a red marker and cell nuclei which were not detected by the method are shown using a yellow marker. While edgeProp is able to separate clusters of nuclei, they are limited by the strength of the edges, resulting in several nuclei which are not detected [De Vylder 13a].

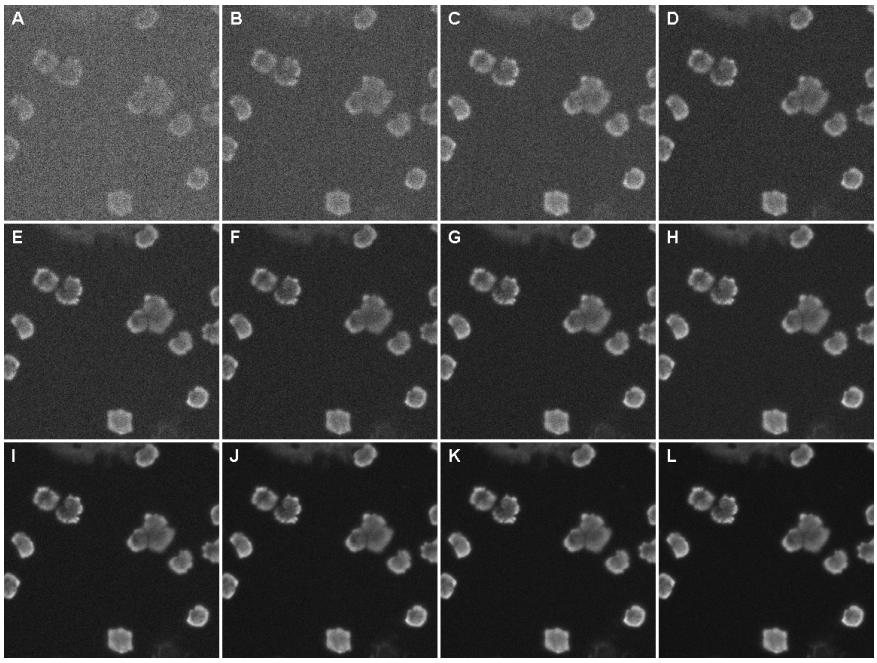


Figure 6.23: An example of the influence of capturing time on a microscopic image. These images are captured using respectively 5, 10, 20, 30, 40, 50, 75, 100, 200, 300, 400 and 500 ms [De Vylder 13a].

the range of diameters of desirable segments. The state-of-the-art methods significantly decrease in detection accuracy if the micrographs are more densely packed with cell nuclei (Table 6.7). From the state-of-the-art techniques the edge based approach is generally more robust against touching nuclei, but still performs significantly less than the proposed method. The method presented in this work not only performs best for relatively sparse spaced nuclei, i.e. 30% of touching nuclei, but also for densely clustered nuclei. The fluorescent dye also remains more compact in dense clusters of nuclei, since there is less free space to dissolve the dye.

In a third experiment we test the robustness against low exposure times. The influence of the capturing time on the quality of the image is illustrated by an example in Fig. 6.23. For visual inspection by an expert, an exposure time of 200 ms is advisable. This results in good contrast and low noise images. However for automatic detection lower exposure times also give adequate results, resulting in faster scanning times, lower influence of phototoxicity and less risk on photobleaching [Stelzer 98, Keller 08]. The ideal exposure time depends on the required accuracy and on the application, e.g. in vivo experiments will require lower scan-

	data set 1			data set 2			data set 3		
method	P	R	F	P	R	F	P	R	F
EdgeProp	0.823	0.895	0.858	0.774	0.847	0.809	0.649	0.801	0.730
CellC	**0.956**	0.824	0.885	0.880	0.694	0.776	0.717	0.488	0.581
CellProfiler	0.855	0.953	0.902	0.794	0.887	0.838	0.500	0.605	0.547
Proposed	0.948	**0.969**	**0.959**	**0.905**	**0.969**	**0.936**	**0.801**	**0.950**	**0.870**

Table 6.7: Comparison of the detection results from different state-of-the-art methods for nuclei of peripheral blood mononuclear cells (P = precision, R = recall, F = F-measure).

	EdgeProp			CellC			CellProfiler			Proposed		
exposure time (ms)	**P**	**R**	**F**	**P**	**R**	**F**	**P**	**R**	**F**	**P**	**R**	**F**
5	0.288	0.560	0.381	0.053	0.028	0.037	0.017	0.291	0.033	**0.476**	**0.679**	**0.560**
10	0.384	0.662	0.486	0.030	0.045	0.036	0.030	0.633	0.057	**0.673**	**0.835**	**0.745**
20	0.467	0.769	0.581	0.256	0.148	0.188	0.097	0.744	0.171	**0.790**	**0.899**	**0.841**
30	0.529	0.778	0.630	0407	0.297	0.343	0.196	0.778	0.313	**0.815**	**0.912**	**0.861**
40	0.524	0.778	0.626	0.548	0.387	0.454	0.277	0.829	0.402	**0.799**	**0.906**	**0.849**
50	0.578	0.793	0.669	0.616	0.393	0.480	0.334	0.833	0.477	**0.826**	**0.931**	**0.875**
75	0.576	0.774	0.661	0.810	0.440	0.570	0.416	0.829	0.554	**0.831**	**0.925**	**0.875**
100	0.577	0.759	0.656	**0.861**	0.432	0.576	0.451	0.840	0.587	0.836	**0.932**	**0.882**
200	0.634	0.761	0.692	0.853	0.424	0.567	0.496	0.855	0.628	**0.872**	**0.923**	**0.897**
300	0.629	0.743	0.681	**0.880**	0.455	0.600	0.496	0.829	0.620	0.874	**0.923**	**0.898**
400	0.643	0.741	0.688	0.888	0.464	0.610	0.532	0.870	0.660	**0.892**	**0.931**	**0.911**
500	0.640	0.733	0.684	0.875	0.461	0.603	0.519	0.853	0.645	**0.889**	**0.931**	**0.909**
1000	0.670	0.729	0.700	0.887	0.485	0.627	0.545	0.868	0.670	**0.896**	**0.921**	**0.908**

Table 6.8: Comparison of the detection results from different state-of-the-art methods for nuclei of peripheral blood mononuclear cells captured using different exposure times (P = precision, R = recall, F = F-measure).

ning times due to photo-toxicity. State-of-the-art methods clearly lack robustness against low exposure times (Table 6.8), while the proposed method shows accurate detection results, even for low exposure times. For an exposure time of 30 ms, 15% of the 200 ms used for human operators, the proposed method still results in an F-measure of 0.861. All processing work was performed on an Intel i7 Q720 1.6 GHz CPU processor with 4 GB memory.[2] It took an average of 26.944s for the analysis of a single image from the BBBC004v1 data set from the Broad Bioimage Benchmark Collection, i.e. for the analysis of an image of 950×950 pixels. While this method is computationally more demanding than state-of-the-art methods (cellC, cellprofiler and edgeProp require respectively 4.33s, 5.80s and 9.78s), the difference does not prevent practical use. Furthermore the analysis of a data set can be offloaded to a dedicated server. By applying such a pipeline architecture, the bottleneck from the image processing work flow is moved from the acquisition, using expensive microscopes, to the analysis on an inexpensive computer.

6.3 Conclusion

In this chapter a number of algorithms for the segmentation of cell nuclei have been discussed. The two main approaches, edge based and intensity based, have been explained using different use cases. Furthermore four new segmentation methods have been proposed.

We proposed an edge based algorithm for the segmentation of individual nuclei in images where nuclei cluster together. In contrast to state-of-the-art methods, the proposed algorithm gives accurate results even if the surrounding contours don't show clear concavities. The proposed method is calculated using an efficient dual scan line algorithm, making it significantly faster than current state-of-the-art, and only requires to keep two rows/slices in memory. This makes the proposed technique interesting for the analysis of big data sets, such as 3D data sets, or image sequences coming from high throughput systems. The method was tested both on synthetic as on real data sets, resulting in good precision and recall. Precision might be further improved if the found segments would be merged using some form of shape constraints.

A second method was discussed to overcome local changes in light intensity. The proposed method is an intensity based method, which makes it robust against blur. The main novelty is that it compares the intensity of a potential segment with the average intensity in the neighbourhood. The method shows accurate results both in bright and dark regions, and in contrast to other "local" methods, such as local thresholding, this without detecting erroneous segments in regions where no nuclei

[2]The complete algorithm is implemented in Matlab R20007b1.

are present. Implemented using a simple linear filter, the method is computational efficient, yet results in accurate nuclei detections. This makes it a good candidate for high throughput analysis of large 3D data sets.

Next, we discussed a novel method for the segmentation of densely clustered nuclei in 3D fluorescence microscopy images. Initially all slices are individually segmented, only for more challenging nuclei configurations, information of neighbouring slices is incorporated. These complex 2D configurations with insufficient confidence are detected using dedicated shape measurements. The method was validated both on simulated synthetic data and based on statistical measurements of the segments. Both for the synthetic and real data, the method outperforms state-of-the-art methods from literature.

Finally, we proposed a novel segmentation technique for 2D micrographs containing dense clusters of circle shaped nuclei. The method finds individual nuclei by jointly optimizing the likelihood of the intensity and the expected shape of the nucleus. This is achieved using the active contour framework, where the proposed method is the result of optimizing a convex energy function. It uses prior knowledge about the shape of cell nuclei, which is done by representing the segmentation result using a dictionary. As far as the authors are aware, this is the first use of shape dictionaries within the context of segmentation. The proposed method was tested both on a benchmark data set and on real microscopic data sets of nuclei belonging to peripheral blood mononuclear cells, showing generic value for the detection of nuclei with an approximately circular shape. The method results in accurate nuclei detection and outperforms state-of-the-art methods, both for precision, recall, F-measure and Dice coefficient. The results show that the method is highly robust against dense cell nuclei clusters and can be used for noisy images captured using low exposure times.

This work is published in a number of international conference proceedings [De Vylder 10f, De Vylder 10g, De Vylder 11e, De Vylder 12a, De Vylder 12e] and in a peer reviewed journal [De Vylder 13a].

Part III

Shape & deformation

7

Cell nuclei mapping

The spatiotemporal dynamics of protein complexes and genome loci, which are located within the nucleus, are functionally linked to cellular health status. In order to study the inherent motion of subnuclear particles, it is essential to remove any superimposed component stemming from displacement and deformation of the nucleus. The actual protein and loci displacements can then be determined by comparing the predicted location with the actual location of the nuclear loci of interest. In this chapter, we discuss different strategies for mapping the nuclear interior, i.e. we calculate for each point in the nucleus where it should move to under the influence of the overall motion and deformation of the nucleus. These mapping procedures enable an accurate estimation of telomere mobility in living human cells undergoing nuclear deformations. Given the large variety of pathologies and cellular processes that are associated with strong nuclear shape changes, contour mapping algorithms have generic value for improving the accuracy of mobility measurements of genome loci as well as intranuclear macromolecule complexes.

Within the limited space of the human cell nucleus, the entire genome as well as many proteins are accommodated in a non-random and time-dependent manner. Of particular interest are telomeres, the very ends of chromosomes, which show a spatiotemporal behaviour that is functionally linked to cellular and organismal condition. Telomeres are arranged in distinct patterns and display mobility regimes at different time scales [Molenaar 03, De Vos 10a]. Diverse biological processes such as telomere maintenance, senescence (i.e. biological ageing) and apoptosis

(i.e. programmed cell death) are associated with characteristic telomere redistributions [Raz 06, Raz 08]. Likewise, altered telomere dynamics are associated with specific diseases such as laminopathies and cancer [Louis 05, De Vos 10a]. Therefore, quantitative studies of telomere dynamics - and by extension nuclear dynamics - may help reveal novel mechanisms of dysfunction and disease. However, mobility analysis in four-dimensional image data sets is seriously hampered by global cell motion and nuclear deformation. This superimposes a global motion on the submicron-scale dynamics, which in turn creates a significant bias on the accuracy of the measurements. It is therefore imperative to remove the superfluous motion by means of image registration.

A wide range of registration techniques have been proposed to remove 2D and 3D motion of cell nuclei using rigid transformations, i.e. translation and rotation [Gerlich 01, Rieger 04, Wilson 06]. However, the unwanted global and local expansions or contractions, cannot be adequately modelled by rigid transformations. To cater for complete non-rigid transformations, landmark-based methods have been proposed [Zitova 03, Mattes 06, Matula 06], but the results of these methods are completely dependent on the choice of robust landmarks, which is not always straightforward. Furthermore, these landmarks should only be affected by global motion induced by nuclear displacement and deformation, and not by the local telomere motion which has to be analysed. Registration methods based on the overall intensity of the nucleus have been proposed [Zitova 03, Myronenko 10, Kim 11], but they have the risk of biasing the results. Using intensity, including intensity of the telomeres can result in over-fitting the telomere displacement mapping. This displacement can range from small Brownian motion to large telomere displacements, as can be seen in Fig. 7.1, where a telomere moves over 10 μm in 26 min.

Recently a set of dedicated registration algorithms have been developed for the mapping of cell nuclei. These methods avoid the hazards of intensity or landmark based registration, but instead focus on the deformation of the nuclear contour. These methods first start by segmenting the nucleus. Using this segmentation, a proper non-rigid displacement mapping is calculated.

This chapter is arranged as follows: in the next section we briefly discuss the algorithm that can be used for segmenting isolated nuclei. Section 7.2 describes the state-of-the-art method for removing all global motion from healthy cells in interphase, the part of the cell cycle in which the cell spends most of its time. In section 7.3 we propose a our own novel mapping strategy that can be used for pathological nuclei. This proposed method is not constrained to ellipsoid-like nuclei and as such can be used for the mapping of aberrant nuclei corresponding to pathological cells. Section 7.4 recapitulates and concludes.

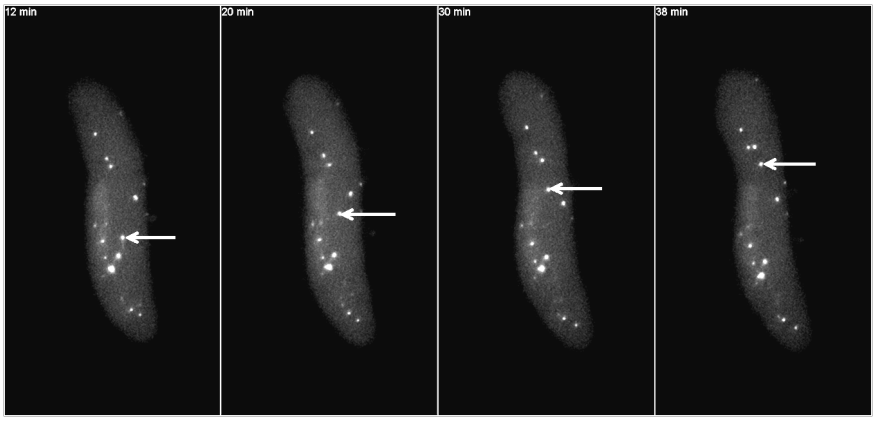


Figure 7.1: An example of telomere mobility. Within a time frame of 26 minutes. the highlighted telomere moved approximately 10 μm, while the cell nucleus remains stationary.

7.1 Segmentation

Since contour based registration methods estimate the nuclear geometrical changes based on the deformation of the nucleus, the nucleus needs to be segmented first (Fig. 7.3.(a)). For this application we assume to study isolated nuclei. This makes the segmentation task much easier, and allows us to use simple, common techniques, omitting more complex methods such as those described in chapter 6. First the image is preprocessed using a low-pass filter in order to suppress noise and non-uniform staining. The resulting image is then binarized using Otsu's thresholding method, which calculates a threshold by minimizing the variance of both the foreground and background [Otsu 79, Sezgin 04]. Due to subnuclear particles such as nucleoli and remaining noise, not all nucleus pixels are correctly classified as corresponding to the nucleus and also some background pixels will wrongfully be considered as foreground. This is corrected with a morphological opening and closing (see chapter 4). Finally, the contour delineating the segment is parametrised based on the arclength and filtered with a mean filter in order to obtain a smooth contour. This is important for the contour matching which we will discuss in the next two sections. For a more detailed discussion on cell nuclei segmentation, we refer to chapter 6.

7.2 Interphase mapping

To handle both global and local expansions or contractions, a mapping based on a spherical model has been proposed [Gladilin 08]. This model captures most of

the cellular and nuclear displacements that occur in normal cells, i.e. healthy cells during interphase. In a first step, an ellipsoid shaped mesh template is aligned with the 3D segment representing the nucleus. This alignment is rather straightforward, because the segment is close to elliptical. The template is translated in such a way that its centroid coincides with the centre of mass of the segment. Then the mesh is scaled and rotated such that its principal axis are identical to the principal axis of the ellipsoid that has the same second order moments as the segment (see chapter 8 for more details about shape moments).

In a second phase the template mesh is non-rigidly aligned. Given the ellipsoidal shape of a healthy nucleus, we can assume this local deformation to be small. Hence a "propagation by normals" technique similar to the optimization in section 6.2.2.3 can be used, i.e. the 3D mesh node $\mathbf{u}$, is displaced along the surface normal $\mathbf{n}$:

$$\mathbf{d} = \mathbf{u} + \lambda \mathbf{n}, \tag{7.2.1}$$

with the point on the segments boundary corresponding with $\mathbf{d}$ and λ the Euclidean distance between $\mathbf{u}$ and $\mathbf{d}$. Based on elasticity theory and finite element methods, these boundary conditions can be extended to a displacement map for the full nucleus [Gladilin 07, Gladilin 08].

7.3 Nuclear mapping of pathological cells

The previous method gives good results for healthy cell nuclei during interphase. However, under certain pathological conditions, the nuclear shape alters dramatically [Capo-Chichi 09], creating a situation in which an elliptical model can no longer be applied. A number of examples of aberrant nuclei are shown in Fig. 7.2. Pathologies resulting in aberrant nuclei comprise cancer and laminopathies, a gene mutation related to diseases such as muscular dystrophies, lipid dystrophies and premature ageing syndromes [De Vos 10a]. As can be seen in Fig. 7.2, the ellipsoid-like shape-assumption is clearly violated. This would result in both bad global, and erroneous non-rigid local mapping. Hence, a new method is required that allows reliable motion measurements of subnuclear features in a deformable volume, without imposing any shape constraints. In this section, we propose a novel approach, which maps two nuclei between consecutive time points without imposing constraints (e.g. shape constraints), nor requiring any landmarks. A workflow was created of which the major aspects are described below in more detail and which are schematically summarized in Fig. 7.3. First the image is segmented according to the algorithm described in section 7.1, then based on the contour of the segment a signature is calculated. The signatures of

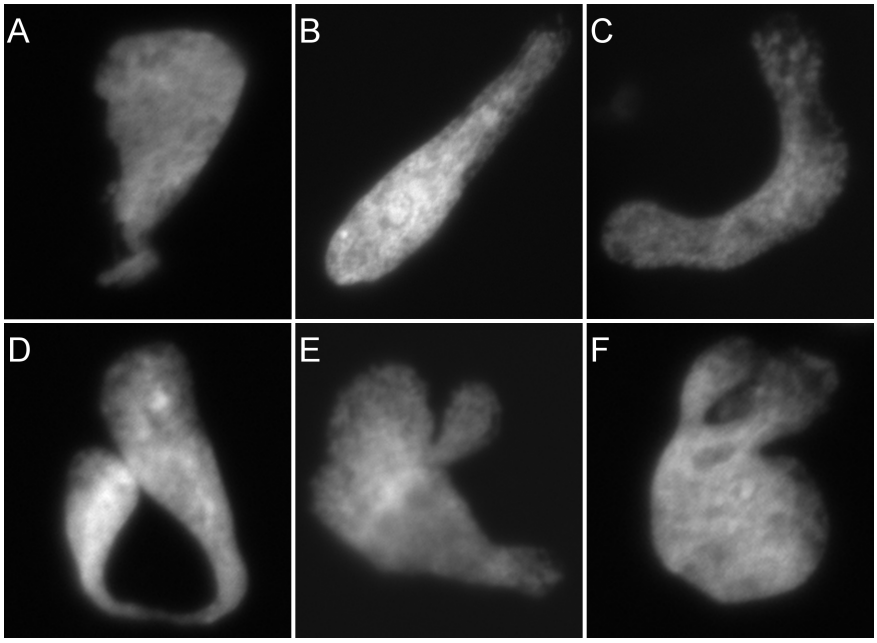


Figure 7.2: A couple of different phenotypes in aberrant nuclei. These aberrations range from random shaped nuclei (A), elongated nuclei (B,C), folding (D), a bulge (E) or crevices (F).

two consecutive time frames are matched, which results in a point-to-point correspondence. Based on this matching, a mapping of the entire nucleus is calculated using polyharmonic splines. This mapping predicts the location of nuclear loci, solely based on the deformation of the nucleus. The algorithm was tested both on simulated images and on recordings of human cells undergoing artificially induced as well as naturally occurring deformation. In the following subsection we will elaborate on each of these steps.

7.3.1 Contour matching

Using the contours of the found segments (section 7.1), we find the correspondence between contours delineating the nucleus in consecutive time frames: points belonging to the contour in one frame will be linked to the corresponding points on the contour in the next frame. To this end, we first calculate a signature $s_i(.)$ of the contour in time frame i. This is a function describing the contour in one dimension. This signature is then used for the actual matching. The contour is described as $\mathbf{c}(t)$=$(x(t), y(t))$, with $t \in \mathbb{Z}$ and $\mathbf{c}(n+t) = \mathbf{c}(t)$, i.e. the contour is periodical. A wide variety of shape describing signatures have been proposed in literature [Kindratenko 03, Zhang 04], i.e. a one dimensional function based on $\mathbf{c}(t)$. Within this work, the following four different signatures were tested (Fig. 7.3.(b)):

- A complex valued signature: $s(t) = x(t) + j\,y(t)$, where $j^2 = -1$.
- A normalized centroid distance signature:

$$s(t) = \frac{\sqrt{(x(t)-x_c)^2 + (y(t)-y_c)^2}}{\sum_{i=1}^{n} \sqrt{(x(i)-x_c)^2 + (y(i)-y_c)^2}}$$

  , where (x_c, y_c) is the centroid of $\mathbf{c}(t)$, i.e. $x_c = \frac{1}{n}\sum_{i=1}^{n} x(i)$ and $y_c = \frac{1}{n}\sum_{i=1}^{n} y(i)$. Note that this signature is normalized such that $\sum s(t)\,dt = 1$. By doing so, this signature is scaling invariant, i.e. invariant to isotropic expansion and contraction of the nucleus.
- A directional signature: $s(t) = \cos^{-1}\left(\frac{y(t+1)-y(t-1)}{\|\mathbf{c}(t+1)-\mathbf{c}(t+1)\|}\right)$. This signature corresponds to the orientation of the tangent line at point t with regard to the x-axis.
- A curvature signature: $\|\frac{d^2\mathbf{c}(t)}{dt^2}\|$. The curvature signature has low values at places where the contour is smooth and high values at places where the contour has sharp angles.

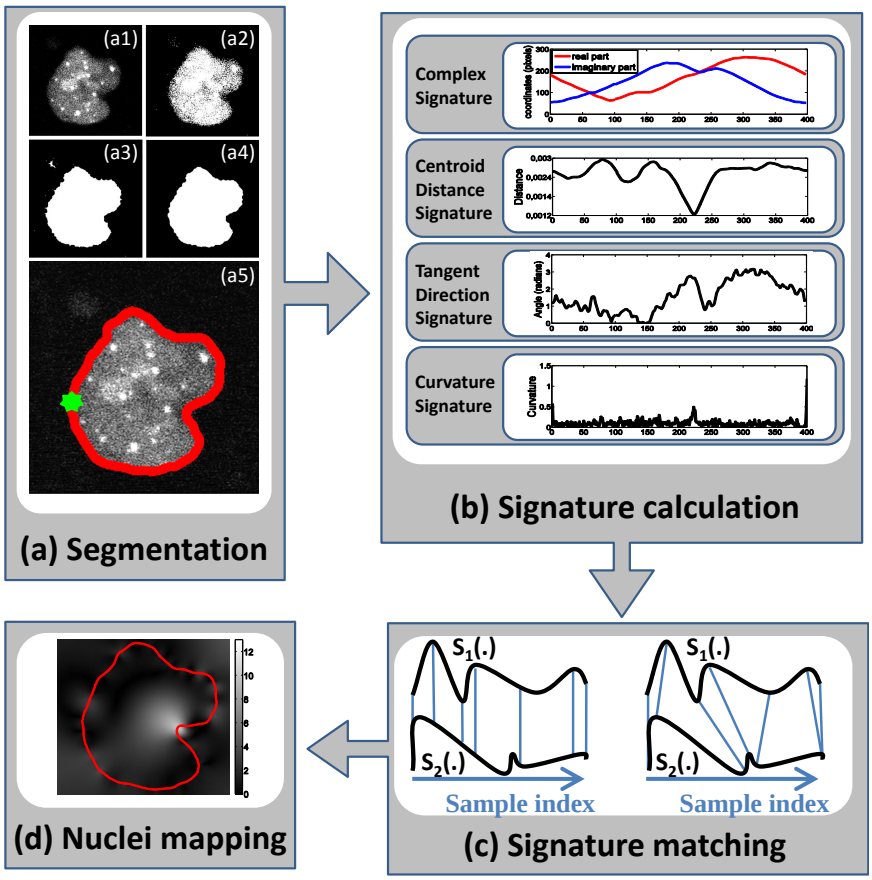


Figure 7.3: Major steps in the workflow of the contour matching registration method: (a) segmentation of the nucleus for the generation of nuclear outlines (red) on the raw image (a1) by consecutive low pass filtering and Otsu thresholding (a2), morphological closing (a3), morphological opening (a4), detection and low-pass filtering of the contour (red outline, a5); The green star indicates the starting point of the contour (b). The contours are transformed into one-dimensional contour signatures. Four different signatures were tested: a complex signature, a normalized centroid distance signature, a tangent direction signature and a curvature signature; (c) The signatures, s1(.) and s2(.), of two consecutive frames are matched according to the dynamic time warping technique. Blue lines represent matched samples. On the left, a naive approach is shown where points between the two signatures are simply matched consecutively ($\mathbf{s}_1(t) = \mathbf{s}_2(t)$). On the right, matching is done with the minimal cumulative difference, or DTW distance. From this contour matching the warping function is derived; (d) Based on the contour matching between consecutive frames, a full nucleus mapping is calculated using polyharmonic spline interpolation. Gray values represent the magnitude of the displacement vector (in pixels) [De Vylder 11c].

Although this is the first use of signature functions for global registration, they already prove their merit for many other applications. For a more detailed overview of signature functions we refer to [Van Otterloo 88, Wood 96, Trier 96] [Kindratenko 03, Zhang 04].

The matching criterion we applied for two signatures is the Dynamic Time Warping (DTW). DTW is also known as the dog-man distance, in analogy to a man walking a dog on a leash. Both walk different paths (i.e. the two signatures) with the potential to walk at different pace depending on the length of the leash (i.e. the warping function), but with the ultimate goal of maintaining overall closest proximity between man and dog (i.e. the DTW distance). In mathematical terms, the DTW distance is the smallest cumulative difference between both signatures $s_1(.)$ and $s_2(.)$, which is found by changing the scanning speed/sampling density (the length of the leash) along the contour signature (Fig. 7.3.(c)):

$$d_{DTW}(s_1(.),\, s_2(.)) = \min_{\gamma(.)} \sum_t \|s_1(t) - s_2(\gamma(t))\|, \tag{7.3.1}$$

where $\gamma(.)$ is a warping function, i.e. a monotonic function mapping all n elements of $s_1(.)$ onto all m elements in $s_2(.)$. The optimal warping function which minimizes the DTW distance will be used for calculating the mapping of the nuclei between consecutive frames. For discrete contour signatures, the DTW distance (and implicitly also the warping function) can be calculated in an efficient way using an $n \times m$ alignment matrix $D(.,.)$, for which:

$$D(i,j) = d(i,j) + \min(\alpha D(i,j-1),\, \alpha D(i-1,j),\, \alpha D(i-1,j-1)) \tag{7.3.2}$$

with $d(i,j) = \|s_1(i) - s_2(j)\|$ and α a real number greater than or equal to one. In order to penalize warping functions where too many points of one signal are matched to a single point in the other signal, we have set $\alpha = 1.25$. This was empirically determined to be a good value for matching nuclei contours in our datasets. Starting from $D(n,m)$, i.e. the DTW distance, the warping function is obtained by recursively finding the minimum D value for a matrix element within its direct neighbours in the alignment matrix, down to its starting point.

The DTW distance is uniquely defined between two finite time series. Although signatures are infinitely long, they are periodical (with period n). In other words between 1 and n, the signature can be considered as finite in time. This, however, assumes that the starting points correspond between frames. If this is not the case, a wrong DTW distance and suboptimal warping function will be obtained. This potential pitfall is avoided by calculating the DTW warping function twice: once to find the offset in starting points and a second time to calculate the exact warping function. When starting points do not correspond between signatures of consecutive frames, the starting point in one signature will be matched to k

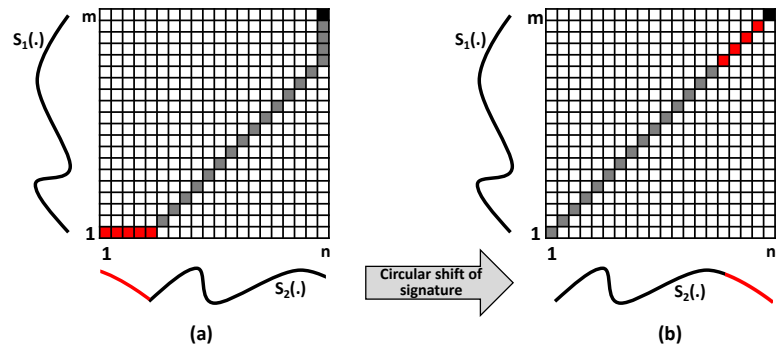


Figure 7.4: Starting points between signatures of consecutive time frames are matched by performing a circular shift. (a) The first DTW procedure generates an alignment matrix, $D(.,.)$, between a signature $s_1(.)$ and a cyclic shifted version of the same signature $s_2(.)$. The optimal warping function c is represented by the path of filled squares. $D(n,m)$ equals the DTW distance (black square). A mismatch in starting points is observed, corresponding to the first k red squares, which are all mapped to the starting point of $s_1(.)$. By circularly shifting these k points, a unique correspondence in starting points will be obtained, resulting in the alignment matrix represented in (b) [De Vylder 11c].

different points in the other signature by the first DTW warping function. A circular shift by k points will create a unique correspondence in starting points. In Fig. 7.4(a), the optimal warping function $\gamma(.)$ is represented by the path of filled squares. $D(n,m)$ equals the DTW distance (black square). A mismatch in starting points is observed, corresponding to the first k red squares, which are all mapped to the starting point of $s_1(.)$. By circularly shifting these k points, a unique correspondence in starting points will be obtained, resulting in the alignment matrix represented in Fig. 7.4(b). Based on this shifted version of the signature the second DTW warping function is calculated.

It may happen that multiple points on one contour match a single point on the second contour, which might lead to unstable solutions in the cell-mapping algorithm (see next subsection). To avoid this situation, both contours are resampled (see Fig. 7.5). Consider two signatures $s_1(.)$ and $s_2(.)$ of consecutive time frames with one point, $s_1(t)$ of the first signature being matched with two points $s_2(a)$ and $s_2(b)$ on the other signature. $s_1(t-1)$ is the last predecessor and $s_1(t+1)$ the first successor of $s_1(t)$ with a unique point-to-point correspondence, $s_2(t-1)$ and $s_2(t+1)$ respectively. The curve between $s_2(t-1)$ and $s_2(t+1)$ will be resampled in such a way that there are as many samples as there are samples between $s_1(t-1)$ and $s_1(t+1)$, making sure that the arclength between the samples $s_1(.)$ is proportional to the arclength of their corresponding samples in $s_2(.)$.

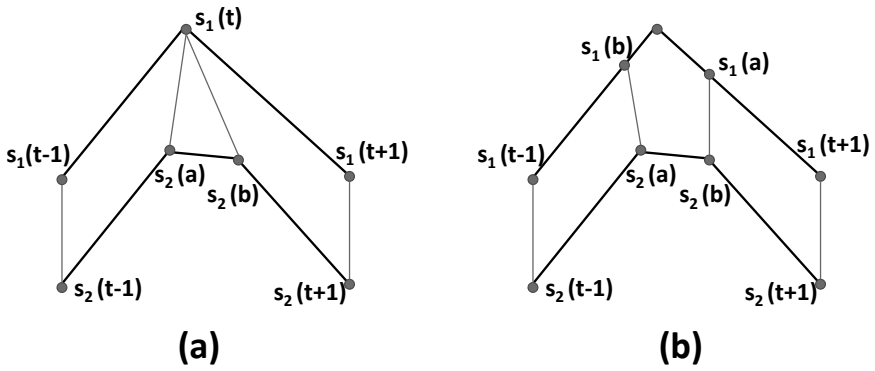


Figure 7.5: The signatures are resampled to obtain a point-to-point correspondence. It may happen that multiple points on one contour match a single point on the second contour. Consider two signatures $s_1(.)$ and $s_2(.)$ of consecutive time frames with one point, $s_1(t)$ of the first signature being matched with two points $s_2(a)$ and $s_2(b)$ on the other signature. $s_1(t-1)$ is the last predecessor and $s_1(t+1)$ the first successor of $s_1(t)$ with a unique point-to-point correspondence, $s_2(t-1)$ and $s_2(t+1)$ respectively. The curve between $s_2(t-1)$ and $s_2(t+1)$ will be resampled in such a way that there are as many samples as there are samples between $s_1(t-1)$ and $s_1(t+1)$, making sure that the arclength between the samples $s_1(.)$ is proportional to the arclength of their corresponding samples in $s_2(.)$. This results in the replacement of $s(t)$ by two new samples, shown as black dots in the figure [De Vylder 11c].

7.3.2 Cell mapping

Based on the contour matching, each point on the contour will now correspond to a unique point on the contour of the subsequent frame. We need to define a function $\psi(\mathbf{p})$, which for each point, $\mathbf{p}$ inside the nucleus in image A, gives the corresponding point $\psi(\mathbf{p})$ in image B inside the nucleus. We will call this the spatial mapping function. We propose to perform an interpolation based on polyharmonic splines:

$$\psi(\mathbf{p}) = a + a_x x + a_y y + \sum_{i=1}^{n} \mathbf{w}_i \phi_k(d(\mathbf{p}, \mathbf{c}(i))), \tag{7.3.3}$$

where the parameters a, a_x, a_y and $\mathbf{w}_j = (w_{j,x}, w_{j,y})$ are a set of weighting coefficients, $\mathbf{c}(i)$ are the contour points, $d(.,.)$ is a distance function and $\phi_k(.)$ is a radial basis function, i.e.

$$\phi_k(r) = \begin{cases} r^k & \text{if } k \text{ is odd,} \\ r^k \ln r & \text{if } k \text{ is even,} \end{cases} \tag{7.3.4}$$

where k is the degree of the basis functions. Note that $k = 2$ is a special case since eq.(7.3.3) then corresponds to the popular *Thin Plate Spline* interpolation

[Bookstein 89]. The weighting coefficients in eq. (7.3.3) are calculated in such a way that:

- contour points are mapped to their corresponding contour points in the other frame, i.e. $\psi(\mathbf{c}(t)) = \mathbf{c}'(t)$, where $\mathbf{c}'(t)$ is the matching contour point of $\mathbf{c}(t)$ based on the optimal warping function;
- the weighting vectors $\mathbf{w}_x$ and $\mathbf{w}_y$ are orthogonal to $\mathbf{x}$ and $\mathbf{y}$, i.e.

$$\sum_{i=1}^{n} w_{i,x} x_i = \sum_{i=1}^{n} w_{i,y} y_i = 0; \tag{7.3.5}$$

- the weighting coefficients sum to zero.

By adding these extra constraints we get $n+3$ linear equations for an equal number of variables. So these constraints provide a mean to easily calculate the weights by solving a linear system, which can be done in a fast way as proposed by Beatson et al. [Beatson 07]. For a more detailed discussion on polyharmonic splines and the influence of the different radial basis functions we refer to [Arad 94, Cofaru 12].

Although polyharmonic interpolation has been proven useful for non-rigid registration, some adjustments are needed in order to use them in this context. The mapping of a point, $\psi(\mathbf{p})$, depends on the Euclidean distance between the point itself and the contour points. However, this distance does not take into account the shape of the nucleus. An example of such a situation is illustrated on the left of Fig. 7.6, where the distance between point A and B corresponds to the length of the dashed line. If the solid grey line corresponds to the nuclear contour, we have to point out that the Euclidean distance does not correspond to the distance that point A would have to move to reach point B. This is because the point can only move within the nucleus, and thus cannot cross the nuclear boundary. As a consequence, the Euclidean distance would be an underestimation of the distance within the nucleus. By ignoring the nuclear shape we erroneously estimate the distance between points which might be close by in Euclidean space, but which can be far away in the nucleus. Although it is difficult to estimate the exact error introduced by using the Euclidean distance, we propose using the geodesic distance instead, i.e. the force that displaces a contour point will induce a displacement within the cell proportional to the distance a particle has to travel to reach the contour point. The geodesic distance represents the length of the shortest path not crossing the cell wall, as is illustrated on the right of Fig. 7.6. The geodesic distance can be calculated using a fast-marching approach [Peyre 08].

Polyharmonic spline interpolation generally results in a smooth mapping with a risk of the maximum displacement falling inside of the nucleus (Fig. 7.7). Assuming that the interior displacement is mostly due to the deformation of the nuclear

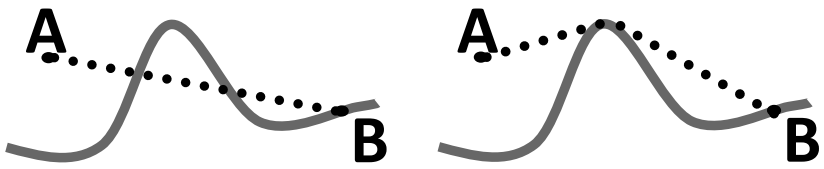


Figure 7.6: An example of two different distance functions: the Euclidean distance on the left, the geodesic distance on the right. The distance between points A and B corresponds to the length of the dashed line, where the solid grey line represents the nucleus contour [De Vylder 11c].

contour, we expect the magnitude of the displacement vector to be maximal along the contour of the mapping and not inside of the nucleus. Especially for non-convex contours there seems to be an overestimation of the motion (see Fig. 7.7.D). To avoid the maximum displacement from falling outside the contour boundary, the derivatives of the mapping normal to the contour were set equal to zero. This was done by duplicating contour points at a short distance (1 pixel) outside of the nucleus before mapping, a technique we refer to as contour padding (Fig. 7.7).

7.3.3 Telomere tracking

In order to study telomere dynamics over time, detected telomeres in consecutive time frames are matched using the Iterative Closest Point method (ICP). This results in a matching where each telomere is matched with exactly one telomere in the next frame. ICP is a greedy algorithm, i.e. the closest telomere points are matched first and this decision is not revised in the remainder of the algorithm. This is iteratively repeated until all telomeres are matched. Other approaches striving for a global optimum according to a predefined criterion could also be used [Kuhn 55, Meijering 06, Zhang 08, Padfield 11]. This greedy approach however has the benefit of simplicity, while achieving adequate results.

7.3.4 Experimental results

First, the mapping algorithm was validated on an artificial sequence using a Monte Carlo simulation. The nucleus was approximated by a flattened ellipsoid ($80 \times 50 \times 40$ voxels) in which 60 random points were generated, representing telomere positions. Then the ellipsoid was deformed using a random affine transformation composed of 3 separate components (of uniform distribution): the three axes of

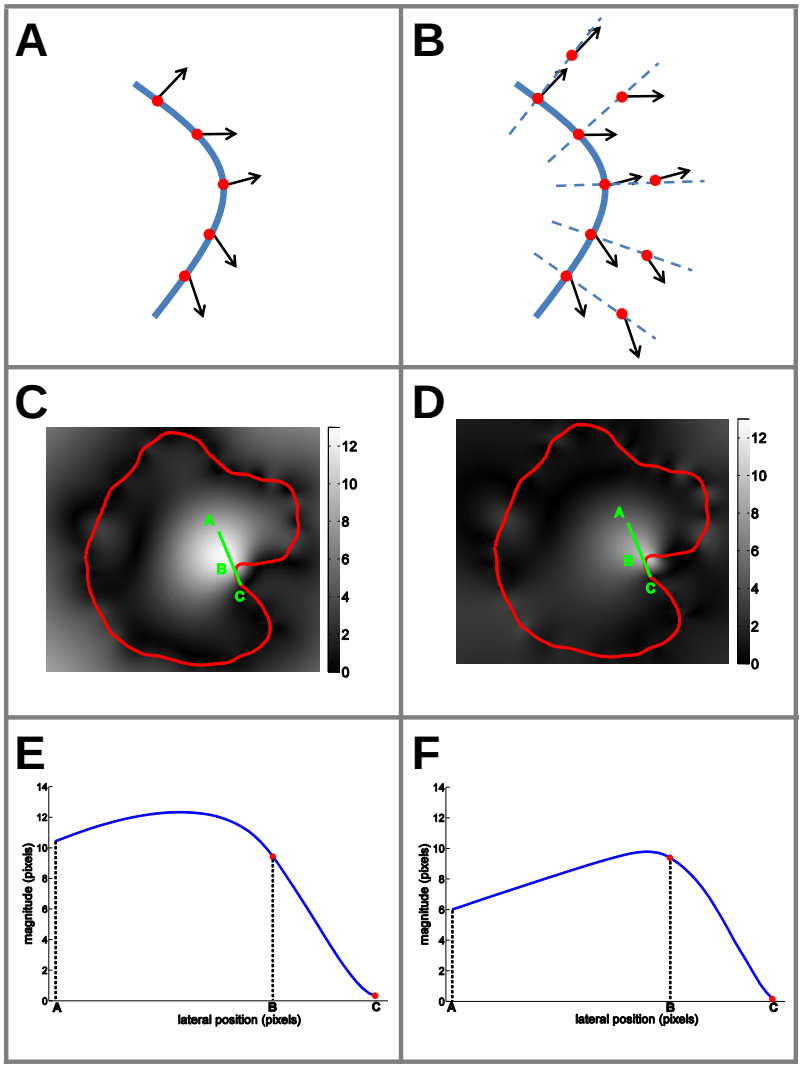


Figure 7.7: Illustration of the effect of contour padding to improve nucleus mapping. (A,B) Contour padding consists of duplicating contour points (red dots) having identical displacement vectors (black arrows) at a short distance outside of the nucleus, normal to the contour, before mapping; (C,D) mapping of the nuclear deformation (based on polyharmonic spline interpolation) between two consecutive time frames before (A) and after (B) contour padding. Gray values represent the magnitude of the displacement vector in pixels. (E,F) Profiles representing the magnitude of the displacement vector along the green line (from lateral position A to C) in (C,D) respectively. As shown in figure (E), without contour padding, the maximum displacement is inside of the nucleus and is approximately one third higher than on the contour (point B). Assuming that the interior displacement is mostly due to the deformation of the nuclear contour, it is unlikely for a displacement inside of the nucleus to exceed that of the boundary. After contour padding (F) the maximal displacement coincides better with the contour [De Vylder 11c].

the ellipsoid were independently rescaled (between a factor 0.9 and 1.1), the ellipsoid was rotated (between $-5\,^\circ$ and $+5\,^\circ$ along the long axis of the lateral plane and between $-10\,^\circ$ and $+10\,^\circ$ along the axial axis) and translated (between -10 and $+10$ pixels in the lateral direction). $10\,000$ simulations were performed resulting in average telomere displacement of 8.7 ± 15.7 pixels. Subsequent contour matching and point tracking resulted on average (all transformations combined) in a remaining telomere motion of 2.4 ± 3.5 pixels (or 72.5% reduction), only resulting from the motion and deformation of the nucleus. This Monte Carlo simulation was repeated with only transformations in the lateral plane, i.e. without the rotation around the ellipsoid's axis and without the scaling in the non-lateral direction. The simulation of only lateral transformations resulted in an average remaining telomere motion of 2.2 ± 3.7 pixels (74% reduction).

As it is not known what the actual mobility of objects in a nucleus will be, it is difficult to validate the algorithm on a real dataset per se. To test the accuracy of the algorithm we used a specially prepared dataset with maximal resemblance to the natural situation (i.e. real nuclei imaged with fluorescence microscopy): ECV-TRF2 cells were acquired during fixation with methanol (Fig. 7.8). From the moment methanol is added, all active and intrinsic motion is blocked. At the same time, the cold methanol induces a contraction of the entire cell, including the nucleus, which persists for about 10 to 20 min. In this setting, telomeres only become displaced by the deformation of the nucleus and the telomeres' centres can therefore serve as ground truth for validation, i.e. the performance of the algorithm can be judged by its relative minimization of the displacement error. Although the induced shrinking is relatively isotropic, not all nuclei deform in the same way, not all parts of the nucleus contract equally and not all telomeres relocate equally far. 97 telomeres originating from 3 different nuclei were manually tracked during 7 frames. This resulted in 582 ground truth telomere matches. The influence of the different signatures as well as the different radial basis functions used by the mapping was tested. These tests were performed using 60 sample points for the contours; more samples did not yield better results (Fig. 7.9).

The proposed mapping was compared with a rigid mapping and with a state-of-the-art non-rigid mapping based on the Demons registration algorithm [Yang 08, Kroon 09]. For the latter we used an implementation without symmetry criterion [Kroon 09]. The following metrics were used for error estimation: remaining telomere motion, standard deviation of the motion and the number of false matches in the tracking algorithm (Table 7.1). A linear basis, i.e. polyharmonic basis of degree 1, performed best for all four signatures. This can be seen in Table 7.1, where all three measurements are minimal for the linear basis (the order of the radial basis function is denoted by $1^{st}, 2^{nd}, 3^{rd}$ and 4^{th}. From the four signatures tested, the "tangent direction signature" resulted in significantly better results, i.e. lower measurements.

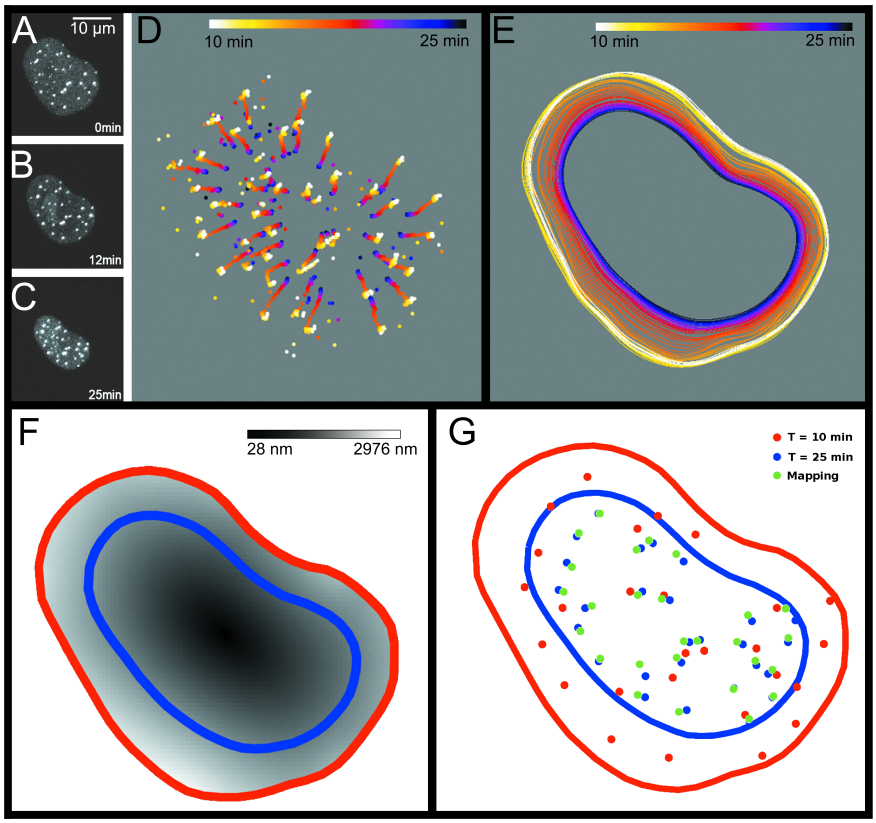


Figure 7.8: Contour mapping applied to a chemically shrunken nucleus. ECV-TRF2 cells were acquired at 30 sec intervals with in fluxo addition of methanol after 5 min of acquisition (A-C) Fluorescent micrographs showing progressive shrinking throughout the 25 min acquisition period (D) Telomere motion over a period of 15 minutes, represented as colour-coded tracks of the telomeric centres of mass. The calibration ranges from 10 min (white) to 25 min (dark blue) (E) Deformation of the nucleus over the same period, represented by the nuclear outlines of consecutive time points (F) The nuclear mapping result after contour matching using a complex signature and a radial basis of degree 1, superimposed onto the outlines of the nucleus at 10 min (red) and 25 min (blue). The gradient displacement calibration bar ranges from 28 nm (black) to 2990 nm (white) (G) The result of telomere mapping after contour matching using a complex signature and a radial basis of degree 1. The red and blue curves are the contours of the nucleus at 10 min and 25 min and the red and blue dots mark the centroids of the telomeres of the respective time points. The green dots represent the location of the telomeres predicted by the mapping, i.e. solely based on the deformation [De Vylder 11c].

			Methanol induced deformation dataset		
			D (nm/min)	σ(nm/min)	FM
Non processed			357.5	26.2	64
Rigid registration			146.6	17.7	31
Demons registration			128.9	12.6	15
Contour matching	Complex	1^{st}	90.5	8.1	6
		2^{nd}	96.0	8.5	7
		3^{rd}	102.8	9.8	7
		4^{th}	107.8	11.5	8
	Centroid distance	1^{st}	99.0	9.5	9
		2^{nd}	110.0	10.4	10
		3^{rd}	116.5	11.5	12
		4^{th}	122.4	13.0	12
	Tangent direction	1^{st}	**64.9**	**6.6**	**3**
		2^{nd}	78.5	8.0	4
		3^{rd}	85.5	9.6	6
		4^{th}	92.0	11.6	6
	Curvature	1^{st}	75.1	7.2	4
		2^{nd}	86.9	8.3	6
		3^{rd}	94.4	10.0	8
		4^{th}	100.4	12.1	9

Table 7.1: Error estimated in registration of an image dataset from methanol treated ECV-TRF2 cells. Three different error metrics were calculated: remaining average telomere displacement D (nm/min), the standard deviation of D and the number of false matches (FM) in the tracking algorithm. These metrics were calculated on the raw dataset (non-processed data) and after registration using rigid registration, a variation on the Demons registration algorithm and the proposed contour-matching algorithm. The latter was tested with four different signatures (a complex signature, a centroid distance signature, a tangent direction signature, and a curvature signature) in combination with polyharmonic spline interpolation with radial basis of four different degrees (1^{st}-4^{th}).

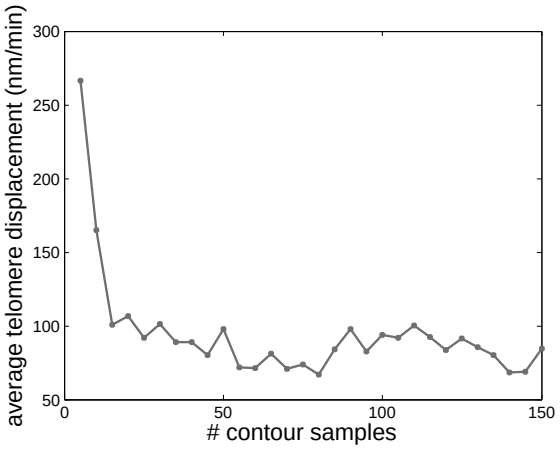


Figure 7.9: The influence of the number of contour samples on the accuracy of the telomere mapping. Calculated on the methanol dataset, using the tangent direction signature and the polyharmonic basis of degree 1 [De Vylder 11c].

Finally, the algorithm was applied to a real dataset from transiently transfected LMNA-/- cells (Appendix A). Nuclei from these cells display dramatic, apparently random deformations, which clearly drive a subset of telomeres in a preferred direction. This makes it difficult to estimate the inherent mobility of the telomeres within the nucleus. 34 telomeres coming from 3 different nuclei where manually tracked over 9 frames. This results in ground truth of 272 telomere displacements. In Fig. 7.10 an example of the mapping is shown. As can be seen, the mapping removes most of the motion stemming from nuclear deformation (Table 7.2). These results show a significant difference compared to non-processed data. Again, these results confirm the biharmonic radial basis and directional signature as the proposed mapping which results in the least telomere displacement. Note that this displacement is not a measure of error since these results come from living nuclei, i.e. there is still telomere motion coming from other biological processes than nuclear deformation. This telomere motion should however be considerably smaller than the motion induced due to nuclear deformation and displacement. So the results obtained with our method are more plausible.

All processing work was performed on an Intel i7 Q720 1.6GHz processor with 4 GB memory.[1] It took on average 2689 ms to calculate the weighting parameters used by the mapping, i.e. $\mathbf{w}_j$, a, a_x and a_y in eq. (7.3.3). For each telomere, the calculation of its mapped location took an additional 103 ms. Approximately

[1] All implementations are done in Matlab R20007b®

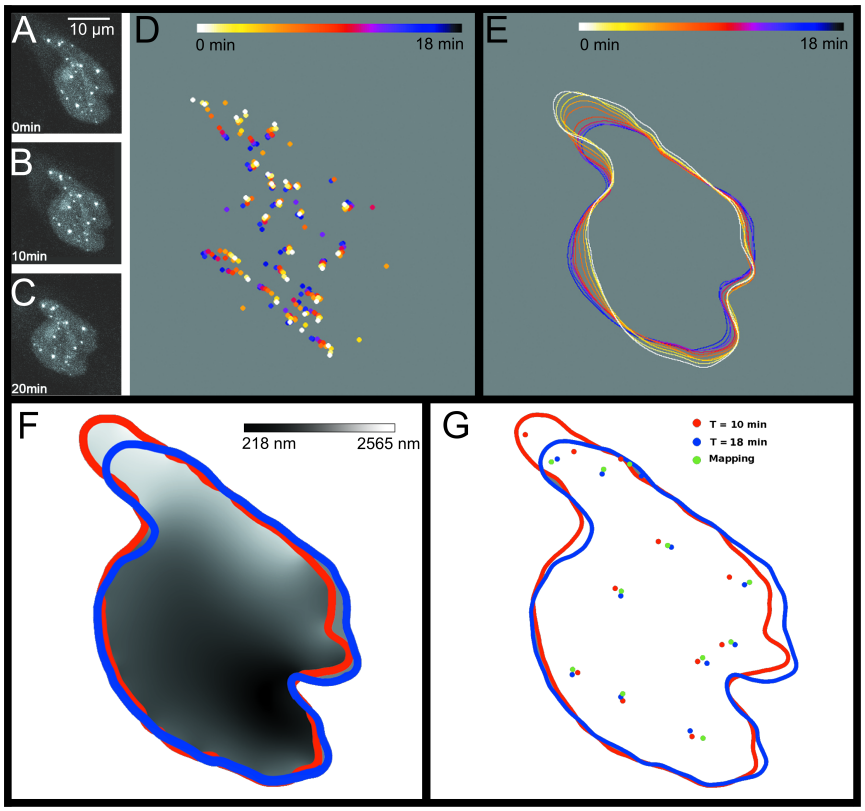


Figure 7.10: Contour mapping applied on naturally deforming nucleus from an LMNA–/– cell, transiently transfected with pmCitrine-hTRF2, acquired at 2 min intervals (A–C) Fluorescent micrographs at different time points. (D) Telomere motion over a period of 18 min, represented as colour-coded tracks of the telomeric centres of mass. (E) Deformation of the nucleus over the same period, represented by the nuclear outlines of consecutive time points (F). The nuclear mapping result after contour matching using a tangent signature and a radial basis of degree 1, superimposed onto the outlines of the nucleus at 0 min (red) and 18 min (blue). The calibration bar represents the relationship between gray values and the magnitude of the gradient displacement (G) The result of telomere mapping after contour matching using a tangent signature and a radial basis of degree 1. Red squares, blue dots, and green triangles represent the positions of tracked telomeres in the first time point (0 min), the last time point (18 min) and the mapping-based prediction, respectively [De Vylder 11c].

			Methanol induced deformation dataset		
			D (nm/min)	σ(nm/min)	FM
Non processed			189.0	41.9	36
Rigid registration			88.0	15.5	10
Demons registration			81.3	12.9	7
Contour matching	Complex	1st	89.1	18.9	15
		2nd	100.6	19.9	17
		3rd	106.5	20.7	22
		4th	109.7	21.3	22
	Centroid distance	1st	75.8	21.3	7
		2nd	71.8	12.3	6
		3rd	75.9	12.9	9
		4th	80.5	13.8	11
	Tangent direction	1st	66.5	**11.2**	**3**
		2nd	**64.7**	11.4	4
		3rd	66.3	11.6	4
		4th	69.3	11.9	4
	Curvature	1st	71.3	13.6	8
		2nd	70.4	13.7	9
		3rd	72.8	14.0	9
		4th	76.1	14.5	11

Table 7.2: Results based on a real dataset from LMNA-/- cells. Three different error metrics were calculated: remaining average telomere displacement D (nm/min), the standard deviation of D and the number of false matches (FM) in the tracking algorithm. These metrics were calculated on the raw dataset (non-processed data) and after registration using rigid registration, a variation on the Demons registration algorithm and the proposed contour-matching algorithm. The latter was tested with four different signatures (a complex signature, a centroid distance signature, a tangent direction signature, and a curvature signature) and polyharmonic spline interpolation with radial basis of four different degrees (1^{st}-4^{th}).

90% of this time was spent on the calculation of the geodesic distance between points. There is no significant difference in computational time depending on the type of signature or on the type of basis function used. This means that for a realistic number of 50 tracked telomeres the total time requirement for contour mapping between consecutive frames is 7839 ms. In comparison, Demons method required 56520 ms for registering consecutive frames. Although the latter method is independent of the number of telomeres it is clearly more time-consuming.

7.3.5 Discussion

Of the four different contour signatures that were tested, the tangent direction signature proved most efficient. Based on this contour matching, the entire nucleus is mapped using polyharmonic interpolation, radial basis of degree one gave the best results.

The primary goal of this registration procedure was to enable an accurate estimation of telomere mobility in living human cells undergoing dramatic nuclear deformations such as human LMNA-/-cells. The nuclei in the 4D CLEM recordings display both nuclear deformation as well as lower level telomere mobility, making it difficult to assess the true performance of the algorithm on these datasets alone. On the other hand, it was difficult to simulate such a multifactorial dynamic system in silico, as this would have required explicit knowledge on the expected biological deformations. Therefore, we have limited the possible deformations in the simulation to affine mappings.

In an attempt to overcome the limitations of simulations, we devised a method for controlled deformation, devoid of additional individual object mobility, with maximal correspondence to the biological topic of interest: a chemically induced shrinking procedure. To our knowledge, this is the first time such an approach was adopted for mimicking deformation of biological material in an uncompromised and controlled manner. Using this approach, we demonstrated a 5.5 fold increase in localization accuracy with respect to the raw data set. Some part of the remaining error of 64.85 nm/min might be due to the contribution of axial components in 3-dimensional movements, which may have been obscured by mapping the deformation in 2D (on projections). However, both cell types under investigation have a flattened nuclear shape in the axial direction suggesting that most of the deformation does indeed occur in the lateral direction [Koopman 06].

Nevertheless, neither axial deformation, nor rolling phenomena can be excluded. In order to estimate the influence of these non-lateral deformations, simulations with all possible 3-dimensional deformations were compared with simulations containing only lateral deformations. This resulted in a net gain of 2% removal of telomere motion compared to the full deformation simulations (74% vs. 72%), which shows that lateral movements indeed cause most of the motion and 2D mapping already removes a major contribution. This is based on the assumption of an affine transformation, whereas nuclei are obviously not restricted to this type of stringent transformations. Future work should therefore focus on extending the method to 3D, e.g. by warping contours on a slice by slice basis, or by using more complex 3D shape matching techniques, before applying a 3D polyharmonic spline mapping. Next to the error induced by 2D mapping, the interpolation and the shape matching will introduce a small error as well as both are approximations

of the real biological deformation.

The general assumption in this work was that the bulk movement of the nucleoplasm could be predicted purely based on the movement of the nuclear boundary. This is supported by particle-tracking microrheology (PTM) studies, which have shown that the intranuclear region is much stiffer than the cytoplasm, and also more elastic than viscous, indicating a strong solid-like behaviour [Tseng 04]. Nevertheless, the nucleus is not a homogeneous structure and consists of transient nuclear microdomains, resulting in local differences in viscosity [Tseng 04]. These differences may explain the remaining localization errors and need also be taken into account when analysing subnuclear particle dynamics in living cells. Therefore, a natural next step in the validation and elaboration of the presented algorithm would be to apply PTM as a model for passive diffusion in the same biological systems [Wirtz 09].

In live cell data it became clear that this type of registration is essential for accurate estimation of the telomere displacements. The average displacement was reduced by a factor 2.8 with the proposed mapping: 189 ± 42 nm/min without registration, and 66 ± 11 nm/min using the mapping. This obviously emphasises the importance of the mapping for accurate mobility estimations. The mapping does not only influence the quantitative measurements of telomere dynamics, but also has a positive impact on the tracking algorithm. The number of false matches obtained with the iterative closest point algorithm reduces with approximately 61%, which is an improvement of 12% compared to registration using the Demons method [Yang 08, Kroon 09].

7.4 Conclusion

This chapter discussed the problem of mapping cell nuclei. It describes the golden standard from literature and presents a new method able to map aberrant nuclei. We have presented a novel registration technique for mapping strongly deformable nuclei. Given the large variety of pathologies (laminopathies, cancer...) and processes (apoptosis, mitosis) that are associated with strong nuclear deformations, we believe this algorithm has generic value by improving the accuracy of mobility measurements of genome loci as well as nuclear protein bodies. The essential novelty of our work is that, in contrast to most of the existing methods, no specific shape constraints are imposed on the contour mapping [Gladilin 08, Mattes 06]. By doing so, it is possible to map contours with a non-elliptical shape, which lack clear and stable landmarks. Some work has been done on generic non-rigid registration without shape constraints [Zitova 03, Yang 08, Tiehua Du 09], but several of these methods use intensity, which is unstable and highly variable in fluorescent

images. Yang et al. proposed a method based on the mapping of segmented nuclei, but their method has difficulties coping with strong deformations, e.g. not all telomeres are mapped inside the nucleus [Yang 08]. The only work, known to the authors, handling strongly deformed nuclei restricts itself to shape analysis [Rohde 08]. It however, does not result in an actual nuclear mapping, which enables the study of telomere dynamics.

This work is published in a number of international conference proceedings [De Vylder 10c, De Vylder 10a, De Vylder 10b, De Vylder 11c, De Vylder 11f, De Vylder 12b] and in a peer reviewed journal [De Vylder 11c].

8

Shape description

Shape is a highly discriminative feature of biological objects that is often used by human observers. Prior knowledge on the shape will help to correctly detect and classify an object of interest. Shape is a natural quantifier for human observers, e.g. a human can intuitively detect similar shapes in a set of random objects often even without any prior knowledge of the objects in the set. For automated computer vision methods, this is far from easy. In order to quantitatively describe shapes, e.g. to enable to the automatic detection of similar shapes, a compact shape representation is often used. A simple example of such a representation could be a small vector of a specific length, the similarity or dissimilarity of multiple shapes could then be expressed by the Euclidean distance between the corresponding vectors. Although a compact shape representation is not necessary, it is certainly interesting for pattern recognition, where it helps to overcome the curse of dimensionality. A good shape representation should preferably be invariant to the location and orientation of the shape, i.e. an object has the same shape, irrespective if it is translated or rotated. Note that for some applications it could be interesting to take the location and orientation into account, but in general this is by convention considered not to be part of the shape definition.

Since shape is an intrinsic property of many biological and biomedical objects, it is a well studied topic for various tasks:

Phenotyping: shape related measurements have been used to quantitatively describe biological specimens [Chitwood 14]. This can be useful to measure

the influence of genetic modifications in plants or the impact of external stress, e.g. draught or wind [Leister 99, Arvidsson 11, Sun 12]. In the context of nuclear phenotyping shape parameters such as curvature statistics or Euler number are useful tools for detecting and describing nuclear pathologies [Yu 09, De Vos 10b, Cattrysse 13], see Fig. 7.2 for an example of different nuclear phenotypes.

Segmentation: in section 3.4 we already elaborated on the benefits of prior knowledge on the intrinsic object properties for the purpose of image segmentation. Obviously, shape is an example of such a powerful intrinsic property and has been exploited in many different robust segmentation algorithms. Prior shape knowledge can range from very general features, such as size or general orientation of a segment [Cremers 08, Klodt 11]; over more specific but still lenient models, e.g. a segment is approximately elliptical [Radke 04, Freedman 05, Schoenemann 07, Charmi 08]; to very restrictive application specific models, e.g. atlas based segmentation where the segment has to belong to a predefined shape space [Cootes 95, Aguado 98, Cootes 01, Aljabar 09, Kauffman 09].

Calculation of population statistics: the previous discussed task aims at getting accurate segmentation results, which is necessary to study an individual object, however, an opportunity arises when studying populations of objects. There it is not essential to correctly segment each individual object, as long as false segmentation results do not influence the measured statistics of the population significantly. This can be achieved by discarding false or doubtful segments, which can be detected using shape features as was shown in [De Vos 10b, Ochoa Donoso 11]. With this approach, one has to take care to avoid bias:, e.g. removing precariously cells that deviate from the more common cell types in a sample. This may lead to bad statistics and the danger is especially large when the deviating cells are the ones of interst. However, with adequate shape descriptors and sufficient understanding of the experiment this approach can result in reliable population statistics.

Shape classification: one of the most straightforward applications is object classification based on shape. An interesting example of this is the "ImageCLEF plant image identification task", where the goal is to map the geographical distribution of specific plant species. This project aims to monitor the biodiversity by analysing pictures of individual leaves. Different shape descriptors have already been investigated for the task of plant identification [McLellan 98, Neto 06, Du 07]. Another relevant group of applications are related to computer aided diagnosis (CAD) systems, e.g. the automatic classification of masses and tumours into benign and malignant masses [Ishiguro 05, Tsui 10, Tahmasbi 11].

Shape retrieval: a similar task to shape classification is shape retrieval. Here the goal is to retrieve similar shapes, given a specific input shape [Milios 00, Zhang 02b, Petrakis 02, Mori 05, Wei 09]. Note that if only one shape could be retrieved, that this results in shape classification. The benefit of retrieving more than one shape is that for difficult cases, a manual user can make the final decision. This can also be used to study deviations in shape, e.g. a microbiologist can search for nuclei with similar shapes in a database with pathological nuclei, thus linking specific nuclear aberrations with specific pathologies.

Shape registration and interpolation: a final application attempts to match the points, or specific parts of a contour to the points/parts of a second contour. An example of such an application was elaborately discussed in chapter 7. Finding the correspondence between contours allows registering two shapes [Gautama 03, Xia 06, Duan 08, Deboeverie 10, Deboeverie 11]. Given the point correspondence it is also possible to interpolate between two shapes [Sethian 99, Jeong 07, Lewin 10]. This can be useful to study the evolution of a shape over time, e.g. in a time lapse sequence, or in space, e.g. the evolution of a 2D contour through a 3D stack.

Due to the many applications, shape representation has been a very active field of research, and has lead to a variety in description methods. In literature, three classes of techniques can be distinguished: simple features, region based methods and the contour based methods. Simple features are a collection of isolated measurements, generally corresponding to simple measurements on the shape, such as the area, circumference, etc. Contour based and region based methods typically result in a discrete function, a vector or a matrix describing the shape. In region based methods, the object is represented by a binary image. Based on this binary image, a shape descriptor vector is calculated. In contour based methods, the points belonging to the contour are first parametrised. This parametrisation is then used to calculate a set of shape descriptors.

Within this chapter we will discuss each of these three approaches in more detail. In the next section we briefly enumerate some of the most popular scalar morphometrics. Section 8.2 gives an overview of region based methods, with an emphasis on different types of moments. Section 8.3 discusses contour based methods: comprising the popular Fourier descriptors and curvature scale space. In this section we also propose a new shape descriptor, that builds upon the framework of Fourier shape descriptors. In section 8.4 the discussed shape descriptors are thoroughly tested for a number of applications: shape reconstruction, shape classification and image content retrieval. Section 8.5 recapitulates and concludes.

8.1 Scalar morphometrics

While shape itself is complex to describe, a lot of simple measurements related to shape can already help in the analysis of biological objects. These measurements generally have the benefit of having an intuitive interpretation, which has made them popular for many biological applications. While these descriptors are often calculated either based on the full region of the object or based on the contour of the object, we mention them here instead of in the next two sections because they correspond to isolated measurements which can be represented by a scalar value. Descriptors which result in a set of measurements, i.e. which can be represented using a vector will be discussed in later sections. We briefly enumerate a number of the most popular scalar shape measurements here:

Size related measurements

- Area/volume of an object in an image is probably one of the most used features for phenotyping [De Vos 10b, Dhondt 12b]. Note that the area is measured in a 2D image, which in many applications is an overhead projection of a 3D object (Fig. 8.1.B). However, for many relatively flat objects, such as A. thaliana rosettes, this is a very reasonable approximation of the 3D area, and as such is correlated to the biomass of the plant [Leister 99, Walter 07].
- The maximal Feret diameter corresponds to the maximal distance between two pixels belonging to an object in the image (Fig. 8.1.C). Used in combination with area, this parameter can be used for the analysis of plants, where it can hint to changes in petiole length.

Elongation related measurements

- Sphericity corresponds to the ratio between the radius of the largest circle that can be enclosed by all contour points over the radius of the smallest circle enclosing the full shape [Du 07].
- Circularity expresses the ratio between the average distance between the contour points and the centroid over the variance on the distance between the contour points and the centroid [Du 07].
- Stockiness is based on the ratio between the area of a segment and its perimeter. It is a useful measurement which for example can be used to detect the serration of leaves [Jansen 09]. Stockiness is defined as $4\pi \times \text{Area}/\text{Perimeter}^2$. This can be seen as a measure of circularity, i.e. a circular object's stockiness is 1. Values vary between 0 and 1.

- Shape eccentricity, not to be confused with the eccentricity of a conic section, corresponds to the ratio between the length of the main inertia axis and the length of the minor inertia axis [Zhang 04, Du 07]

Curvature related measurements

- The average curvature of the contour can help to distinguish between shapes with smooth and "rough" contours [Cattrysse 13].
- Compactness or solidity expresses the ratio between the area of the object and the area enclosed by the convex hull of the object [Jansen 09, Arvidsson 11]. The convex hull of a rosette corresponds to the contour with the shortest perimeter possible which envelopes the rosette. This is illustrated in Fig. 8.1.D, where the compactness corresponds to the white area over the area enclosed by the red contour.
- Previous shape descriptors are rather straightforward and simple measurements, but more complex measurements can also have their value. An example of such a measurement is the internal energy of a an active contour representing the boundary or the centre line of an object (see chapter 4). These energy based terms have been proven useful for the detection of individual nematodes in clusters of nematodes [Ochoa Donoso 10] [Ochoa Donoso 11].

It should be clear that this is not an exhaustive list, but it gives an overall idea of possible scalar morphometrics. Even though several of these shape descriptors are rather ad hoc, these measurements, or a combination of these measurements, have proven there merit for several shape description tasks [Du 07, Jansen 09, De Vos 10b, Chitwood 14]. An example of plant phenotyping using this type of shape descriptors is discussed in sec. 3.5.2.

8.2 Region-based descriptors

These descriptors typically start from a characteristic function or binary image representing the object of interest, i.e. all pixels corresponding to the object are represented by one, whereas all other pixels have a value of zero. A first approach is to transform this image to a wireframe, e.g. by a skeletonisation procedure [Babin 12, Babin 14]. This wireframe can then be used to analyse and describe the shape. This could be done based on statistics of the wireframe, such as the number of bifurcations [Al-Kofahi 08, Dhondt 12b]. A different approach is to use

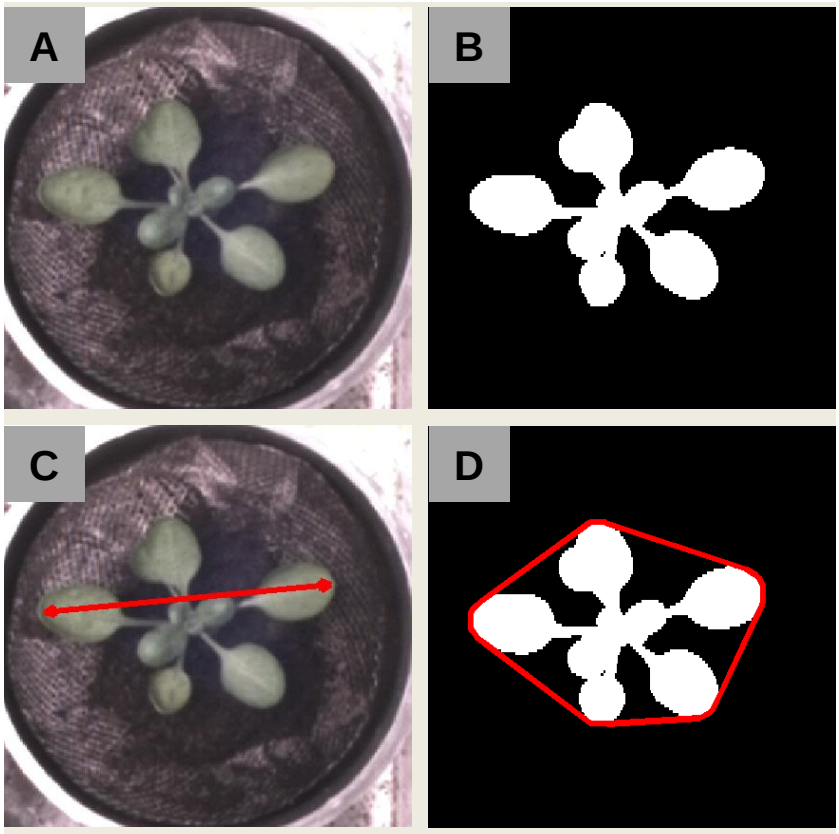


Figure 8.1: An example of different plant morphometrics on a top view VIS image: A) the actual VIS image of a rosette, B) the area detected by Rosette Tracker, C) the Feret diameter of the rosette corresponds to the length of the red line, D) depict the compactness of the rosette, i.e. the ratio of the area corresponding to the actual rosette over the area enclosed by the convex hull, shown as a red line [De Vylder 12c].

this wireframe in a relative way, as a means to compare shapes with each other, e.g. using graph theory [Loncaric 98].

A strategy completely different than the wireframe approach is generalised profiling [Babin 13a]. In this approach a structuring element is defined, e.g. a line or ring, on which a predefined function operates, e.g. a maximum operator. Based on a set of these structuring elements, a set of measurements can be calculated for the binary representation of a shape, i.e. a profile. These measurements can be used as shape descriptors. Examples of shape descriptors that can be calculated using generalised profiling are bounding boxes [Tessens 10], convex hulls, Radon transform [Leavers 00, Kindratenko 03], R-transform [Tabbone 06], etc.

In the remaining of this section we will focus on a different group of methods which describe the image as a linear combination of a set of basis functions.

8.2.1 Geometric moments

The concept of moments in mathematics and physics, i.e. a set of quantitative measures of the shape of a function or distribution, is closely related to describing shape in general. Therefore this concept is adopted in geometric image moments, which have been proposed as a set of shape descriptors [Philips 93, Kindratenko 03, Wei 09, Boveiri 10]. A geometric moment of order $p + q$ is defined as:

$$m_{p,q}(f(.,.)) = \sum_{x=1}^{M} \sum_{y=1}^{N} x^p y^q f(x,y), \tag{8.2.1}$$

where $f(.,.)$ is a binary function with all (x, y) coordinates belonging to the shape correspond to 1 and 0 for all other coordinates and $M \times N$ the dimension of the image. These moments are a measurement of the binary representation of the shape i.e. the $f(.,.)$ function, and not of the shape itself. If a shape is translated, this results in a different $f(.,.)$ function and thus in different moments. This location dependency is solved by centralised moments:

$$\mu_{p,q}(f(.,.)) = \sum_{x=1}^{M} \sum_{y=1}^{N} \Big(x - \frac{m_{1,0}(f(.,.))}{m_{0,0}(f(.,.))}\Big)^p + \Big(y - \frac{m_{0,1}(f(.,.))}{m_{0,0}(f(.,.))}\Big)^q f(x,y). \tag{8.2.2}$$

This has the same effect as translating the binary shape representation, such that the centre of gravity lies at the origin. These centralised moments are still affected by scaling. Scale invariance is achieved by normalizing these centralised moments:

$$\eta_{p,q}(f(.,.)) = \frac{\mu_{p,q}(f(.,.))}{m_{0,0}(f(.,.))^{\alpha}} \tag{8.2.3}$$

with $\alpha = \frac{p+q}{2} + 1$.

8.2.2 Hu-Li moments

A interesting property for shape descriptors is invariance to specific transformations of the shape. Specifically invariance to translation, scaling and rotation are considered essential for general purpose shape descriptors [Xu 08]. The normalized centralised moments from eq. (8.2.3) are translational and scale invariant, but achieving rotation invariance is less trivial. Hu introduced seven rotational invariant moments based on a non-linear combination of geometrical central moments up to the third order [Hu 62]. This set of moments was extended by Li to 52 invariant moments, based on geometrical central moments up to the ninth order [Li 92]. The first twelve invariant moments are shown beneath, note that for readability we have omitted the argument, i.e. the binary $f(.,.)$ function, from the moments.

$$\begin{aligned}
\varphi_1 &= \eta_{20} + \eta_{02} \\
\varphi_2 &= (\eta_{20} - \eta_{02})^2 + 4\eta_{11}^2 \\
\varphi_3 &= (\eta_{30} - 3\eta_{12}) + (3\eta_{21} - \eta_{03})^2 \\
\varphi_4 &= (\eta_{30} - \eta_{12})^2 + (\eta_{21} - \eta_{03})^2 \\
\varphi_5 &= (\eta_{30} - 3\eta_{12})(\eta_{30} + \eta_{12})[(\eta_{30} + \eta_{12})^2 - 3(\eta_{21} + \eta_{03})^2] \\
&\quad +(3\eta_{12} - \eta_{03})(\eta_{21} + \eta_{03})[3(\eta_{30} + \eta_{12})^2 - (\eta_{21} + \eta_{03})^2] \\
\varphi_6 &= (\eta_{20} - \eta_{02})[(\eta_{30} + \eta_{12})^2 - (\eta_{21} + \eta_{03})^2] + 4\eta_{11}(\eta_{30} + \eta_{12})(\eta_{21} + \eta_{03}) \\
\varphi_7 &= (3\eta_{21} - \eta_{03})(\eta_{30} + \eta_{12})[(\eta_{30} + \eta_{12})^2 - 3(\eta_{21} + \eta_{03})^2] \\
&\quad +3(\eta_{21} - \eta_{03})(\eta_{21} + \eta_{03})[3(\eta_{30} + \eta_{12})^2 - (\eta_{21} + \eta_{03})^2] \\
\varphi_8 &= \eta_{40} + \eta_{22} + \eta_{02} \\
\varphi_9 &= (\eta_{40} - \eta_{04})^2 + 4(\eta_{31} - \eta_{13})^2 \\
\varphi_{10} &= (\eta_{40} - 6\eta_{22} + \eta_{04})^2 + 16(\eta_{31} - \eta_{13})^2 \\
\varphi_{11} &= (\eta_{40} - 6\eta_{22} + \eta_{04})^2[(\eta_{40} - \eta_{04})^2 + 4(\eta_{31} - \eta_{13})^2] \\
&\quad +16(\eta_{40} - \eta_{40}) + (\eta_{31} + \eta_{13})(\eta_{31} - \eta_{13}) \\
\varphi_{12} &= (\eta_{40} - 6\eta_{22} + \eta_{04})^2[(\eta_{40} - \eta_{04})^2 + 4(\eta_{31} - \eta_{13})^2] \\
&\quad -16(\eta_{40} - \eta_{40}) + (\eta_{31} + \eta_{13})(\eta_{31} - \eta_{13})
\end{aligned}$$

In [Flusser 93a, Flusser 93b] the concept of non-linear moment combinations is further explored to define a set of four moments which are invariant to affine trans-

formations. This was further improved by introducing a general method for the derivation of any number of affine variant moments [Boveiri 10, Suk 11].

8.2.3 Zernike moments

The geometric moments from eq. (8.2.1) can be generalised by replacing the traditional transform kernel $x^p y^q$ with other kernel functions. The information redundancy of the $x^p y^q$ kernel, due to the non orthogonality of the kernel functions, makes the geometric moments difficult to use for shape classification [Teague 80]. Literature mentions a variety of orthogonal kernel functions for the construction of moments: Legendre polynomials [Shen 96], Chebyshev polynomials [Mukundan 00, Bayraktar 06], Zernike polynomials [Zernike 34,Teague 80], pseudo-Zernike polynomials [Chong 03a], Fourier bases [Zhang 02a, Zhang 02b], etc. Moments based on Zernike polynomials are considered as the gold standard [Kim 00, Chong 03b, Liu 07, Boveiri 10, Tahmasbi 11] for region based shape description and are proven to outperform other (orthogonal) moments for the applications of shape reconstruction [Zhang 04] and shape retrieval [Zhang 03]. Zernike polynomials are orthogonal on the unit disk and are often defined in polar coordinate space. The Zernike polynomial of order p and repetition q is defined as:

$$z_{p,q}(\rho,\theta) = r_{p,q}(\rho)\, e^{jp\theta} \qquad \text{s.t.} \begin{cases} p- \mid q \mid & \text{is even} \\ p \geqslant \mid q \mid \end{cases}, \tag{8.2.4}$$

where $j^2 = -1$, and $r_{p,q}(\rho)$ defined as:

$$r_{p,q}(\rho) = \sum_{s=0}^{(p-|q|)/2} \frac{(-1)^s (p-s)!}{s!((p+ \mid q \mid)/2 - s)!((p- \mid q \mid)/2 - s)!} \rho^{p-2s}. \tag{8.2.5}$$

Using these polynomials results in the following Zernike moments, $\psi_{p,q}$:

$$\psi_{p,q}(f(.,.)) = \frac{p+1}{\pi} \sum_{\rho} \sum_{\theta} z^*_{p,q}(\rho,\theta)\, f(\rho,\theta) \qquad \text{for } \rho \leq 1. \tag{8.2.6}$$

Fig. 8.2 shows the real and imaginary parts of the Zernike polynomials up to order 3. Note that $z_{p,q}(\rho,\theta) = z^*_{p,(-q)}(\rho,\theta)$, where $z^*_{p,(-q)}(\rho,\theta)$ represents the complex conjugate of the Zernike polynomial $z_{p,(-q)}(\rho,\theta)$. One of the benefits of Zernike moments; aside of their usefulness for pattern recognition [Mei 09, Tahmasbi 11]; is that it is possible to reconstruct the original image using the inverse transform:

$$\hat{f}_{p_{max}}(\rho,\theta) = \sum_{p=0}^{p_{max}} \sum_{q} \psi_{p,q}(f(.,.))\, z_{p,q}(\rho,\theta). \tag{8.2.7}$$

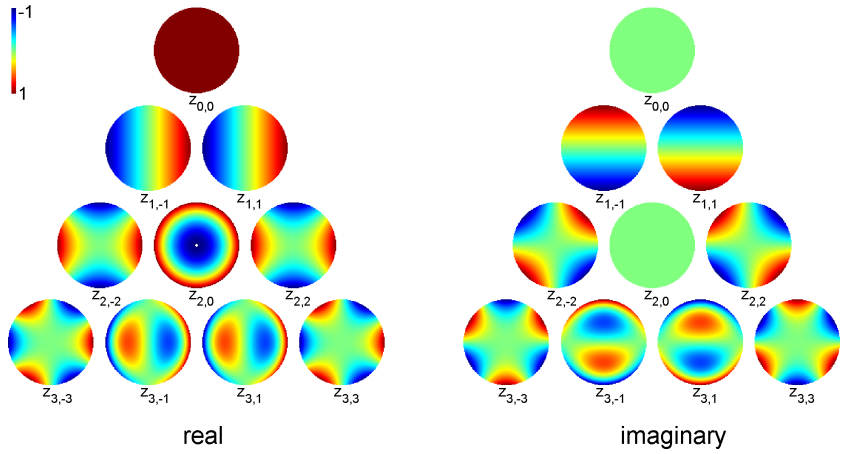


Figure 8.2: The real and imaginary parts of the Zernike polynomials up to order 3.

In theory, the reconstructed image will be perfect if p_{max} is infinite. However a reasonable approximation can be achieved using only moments of lower order. Note that while our shape is represented using a binary function, $f(.,.)$, there is no constraint on the inverse transform that enforces $\hat{f}_{p_{max}}$ to be a binary image as well. This is typically resolved by thresholding the reconstructed image. Figure 8.3 illustrates the reconstruction result using a different number of order of Zernike moments, p_{max}. Figures 8.3(A,C,E,G,I,k) show the actual reconstruction results, which is not binary as already mentioned. Figures 8.3(B,D,F,H,J,L) show the binarised result, achieved by thresholding the respective reconstructions. Fig. 8.3(A,B) is the reconstruction result using moments up to order 5, Fig. 8.3(E,F) used moments up to order 13, Fig. 8.3(I,J) used moments up to order 21, Fig. 8.3(C,D) used moments up to order 29, Fig. 8.3(G,H) used moments up to order 37 and Fig. 8.3(K,L) is the result of using moments up to order 45. Note that the number of descriptors is not equal to the order of the descriptors. For each order of moments a different number of repetitions has to be calculated. While some of these repetitions are the complex conjugate of each other, and can thus be discarded, the total number of repetitions still grows rapidly. For example reconstruction using Zernike moments up to order 5, as is used for the reconstruction in Fig. 8.3(A,B), requires at least 12 Zernike moments. Reconstructions using moments up to order 45, as is used in Fig. 8.3(K,L), already require at least 552 Zernike moments;

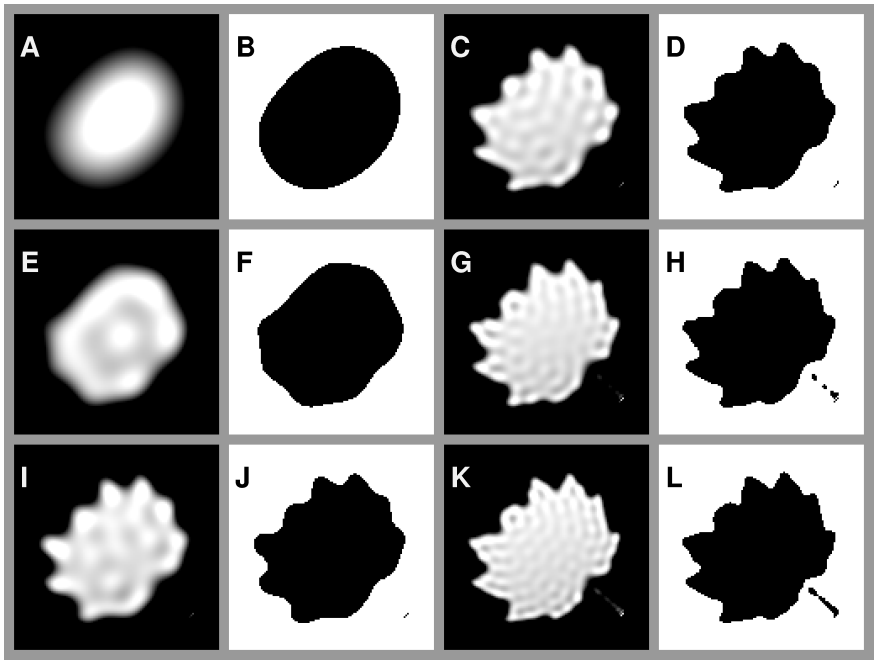


Figure 8.3: An example of shape reconstruction using a different number of order of Zernike moments. The first and third column shows the actual reconstruction, whereas the second and fourth column shows the respective binarised reconstruction. The order of Zernike moments used for the reconstruction are 5 (A-B), 13 (E-F), 21 (I-J), 29 (C-D), 37 (G-H) and 45 (K-L).

8.2.3.1 Invariance

Shape reconstruction itself is a useful property for shape interpolation or as a shape template for segmentation, but for pattern recognition, shape retrieval, phenotyping, etc. it is necessary that a displaced shape results in a set of similar shape descriptors. This can be illustrated by the following example of plant phenotyping: the shape of a plant rosette should result in the same shape descriptors, irrespective of the location of that plant in the image; another example is the automatic classification of segments detected in electron micrographs, structures that correspond to a network of flattened shapes most likely corresponds to endoplasmic reticulum (ER), irrespective if these flattened shapes are oriented horizontally or vertically in the image. So invariance to specific transformations is a desirable property for shape descriptors. Fortunately Zernike moments can achieve invariance to specific transformations using the following steps:

- The effect of a translation is neutralized by translating the binary image representing the image such that the centroid coincides with the origin.
- Invariance to scaling is typically achieved by re-parametrising the image: instead of defining the pixel locations by $\mathbf{x} \in [1, M] \times [1, N]$, each pixel location will be addressed by a different vector. The Zernike polynomials are only orthogonal on the unit circle, so in order to exploit the orthogonality of the basis, the binary image is re-parametrised such that the full object lies within the unit circle. This parametrisation can be achieved by rescaling the object with a factor such that the circumscribed circle of the shape corresponds to the unit circle.
- Rotating the shape with angle α will result in the calculation of the Zernike moments for $f(\rho, \theta + \alpha)$. The real part of the Zernike polynomials is defined solely based on ρ, see eq. 8.2.4, and is thus unaffected by this rotation. Hence, the only effect this rotation has, is a phase shift of $p\alpha$. By using the magnitude of the Zernike moments, the influence of this rotation is discarded, and rotation invariance is achieved.
- Invariance for more general affine transformations can be achieved using more complex approaches, e.g. by ellipse fitting and mapping the ellipses to the unit circle [Amayeh 07, Mei 09].

8.3 Contour-based descriptors

Previous methods work on the whole shape region. However, the contour itself is sufficient to describe a shape. Hence an abundance of methods have aimed

at contour based shape descriptors [Kindratenko 03, Mori 05, Chaker 08, Mei 09, Deboeverie 11, Backes 12, Florindo 12, Hu 13]. The first step in these methods is to represent the contour by a parametric curve. This can either be done explicitly, e.g. by a mapping function $[0, 2\pi] \mapsto \mathbb{R}^2$, or implicitly, e.g. using chain codes or signature functions (see sec. 7.3.1).

This parametric curve can then be used for direct comparison between shapes using specific function similarity metrics [Younes 99, Milios 00, Petrakis 02], or could be used as a first step in the calculation of a set of shape descriptors. In this section we will elaborate on the second approach. This type of methods are absolute instead of relative, i.e. they do not require other shapes to compare with, but can be used on their own. This makes them applicable in a wider range of applications. The next subsection will first discuss methods based on representing the parametric curve using a set of linearly independent orthogonal basis functions. In subsection 8.3.2we represent the parametric curve to a specific type of linearly independent non-orthogonal basis functions.In the last subsection 8.3.3 this is generalised to representing the parametric curve using a frame, i.e. using a set of linearly dependent, non-orthogonal functions.

8.3.1 Shape description using an orthogonal basis

The representation of shape as a linear combination of a set of orthonormal functions has been proven highly efficient for region-based shape description, as described in the previous section. This approach is also widely studied for contour-based shape description. Here not a binary image, but a function, $s(t)$, represents the shape. This function can also be expressed using an orthogonal basis. Just as in region-based shape description different orthogonal basis can be used. Ideally the set of basis functions could be optimised depending on the application. Such trained orthogonal basis comprise eigenshapes [Cootes 95, Cootes 01, Krieger 08, Babaii Rizvandi 08, Kauffman 09], PCA frequencies [Iwata 02, Yoshioka 04, Mei 08], etc.

While a trained basis is optimal for a specific application, trained basis functions also have some drawbacks such as the need for a large representative training-set; sensitivity to geometric transformations; or the sensitivity to the parametrisation, e.g. all shapes have to be parametrised using the same number of sampling points; etc. The dependency on the specific application is also not always desirable, i.e. shape descriptors acquired within a specific application, and thus using a specific orthogonal basis, can not be compared to shape descriptors acquired within a different application, since this will generally use a different basis. Hence the use of more general orthogonal bases has been reported on in literature [Crimmins 82, Bertrand 82, Bartolini 05]. Specially the use of trigonometric functions

seems a popular choice. Consider the trigonometric polynomial of order N:

$$s_N(t) = \sum_{n=-N}^{N} b_n \, e^{jnt} \tag{8.3.1}$$

where b_n minimises the mean squared distance between the reconstruction, $s_N(t)$, and the original shape, $s(t)$. The coefficients b_n are called the Fourier descriptors (FDs). In practice $s(t)$ is a discretised version of a continuous contour. If this discretisation is done properly, i.e. if the Nyquist sampling theorem is met, then perfect reconstruction is possible for N equal to the number of samples. Note that b_n is equal to the n^{th} frequency component of $s(t)$ and can easily be computed using the discrete Fourier transform. This is interesting since low frequencies will correspond to the global shape, whereas Fourier descriptors with a higher index will correspond to high frequency details in the shape. This has as a result that for N smaller than the number of sampling points, the reconstruction approximates the global shape. By increasing the number of Fourier descriptors the shape reconstruction will incorporate more shape details. Fig. 8.4 shows an example of shape reconstruction using a different number of FDs. Fig. 8.4.A shows the original shape whereas Fig. 8.4.B-F shows the reconstruction using respectively 5, 10, 15, 20 and 25 FDs. A low number of FDs results in a very rudimentary approximation of the original shape, whereas increasing the number of FDs gradually adds more detail to the shape, thus improving the approximation.

Note that the FDs are defined for a parametrisation of the shape contour. This is typically done using a complex signature (see section 7.3.1), but the Fourier descriptors can also be calculated on other shape signatures. A thorough comparison of FDs using different parametrisations can be found in [Zhang 05, Elghazal 09, Hu 13]. Most of these parametrisations assume that the shape is described using a closed contour, which is desirable for Fourier descriptors. There are however some parametrisations that have been developed for open contours, these parametrisations introduce periodicity without introducing high frequency components [Weyland 90, Ding 10].

8.3.1.1 Invariance

One of the main advantages of FDs is that for certain transformations, the FDs are either invariant or change in a simple straightforward manner :

- A translation only affects the FD which corresponds to the zero frequency component, i.e. b_0, so except for b_0 the FDs are invariant to translation.
- Scaling the object results in a multiplication of all FDs with a real constant.

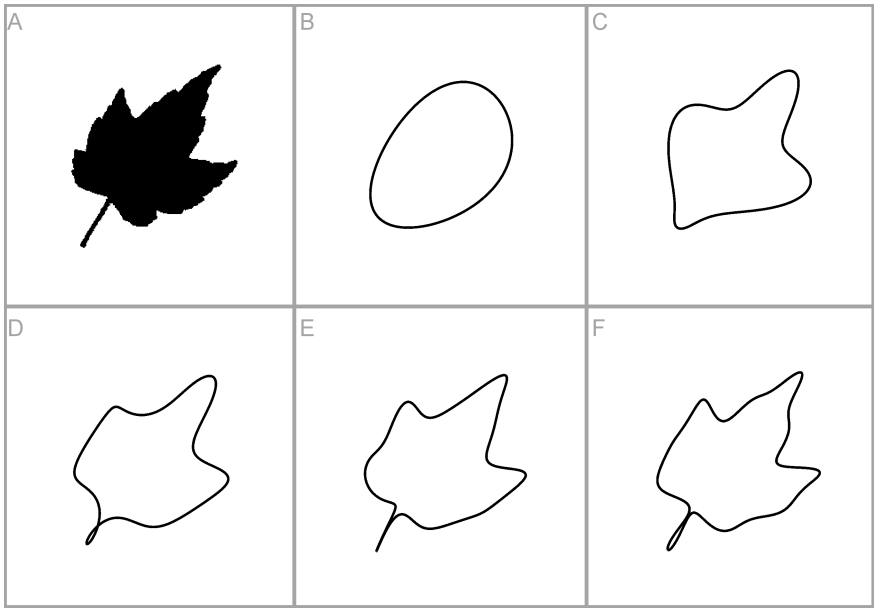


Figure 8.4: An example of shape reconstruction using a different number of Fourier descriptors. A) the original shape B) the reconstruction using 5 FDs, C) the reconstruction using 10 FDs, D) the reconstruction using 15 FDs, E) the reconstruction using 20 FDs and finally F) the reconstruction result using 25 FDs .

- Rotating the object with angle θ results also in a multiplication, but with $e^{j\theta}$.
- Changing the starting point of $s(t)$ results in a phase shift.

These observations make it easy to calculate a set of shape descriptors that are invariant to these shape transformations, i.e. $\frac{|b_{-1}|}{|b_1|}, \frac{|b_2|}{|b_1|}, \frac{|b_{-2}|}{|b_1|}, \frac{|b_3|}{|b_1|}$... The possibility to acquire simple invariant shape descriptors has made FDs popular in a wide range of applications [Aguado 98, McLellan 98, Keyes 99, Folkers 02, Freedman 05, Mahoor 05, Neto 06, Jeong 07, Charmi 08, Chitwood 14]. For invariance to affine transformations, we refer to [Sener 06, Chaker 07, Giannekou 08, Mei 08]

8.3.2 Shape description using a warped orthogonal basis

So far we have discussed the possibility of representing a parametrised curve using an orthogonal basis. While this approach already has proven its merit, the representation of a parametric curve should not be confused with the representation of a shape, i.e. one shape can have many different parametric representations. Consider $C(s(t))$, to be the contour described by $s(t)$, and $C(s(\theta(t)))$, where $\theta(t)$ is a warping function, i.e., any strictly monotonous function that maps the interval $[0, 2\pi]$ onto itself. $C(s(t))$ and $C(s(\theta(t)))$ describe the same contour. This opens the possibility of optimising contour representation by selecting an optimal $\theta(t)$ according to some criterion. The FDs calculate the trigonometric polynomial with the least mean squared distance with $s(t)$, i.e. $d_{MSD}(s(.),\, s_N(.))$. It would be better to calculate the FDs for all warped contours $s(\theta(t))$ and chose the FDs with the smallest error. This results in the following trigonometric polynomial:

$$s'_N(t) = \sum_{n=-N}^{N} c_n\, e^{jnt} \tag{8.3.2}$$

where the coefficients c_n minimise the mean squared distance between $s'_N(t)$ and $s(\theta(t))$ with $\theta(.)$ the most optimal warping function. The G_n coefficients are called improved Fourier descriptors (IFDs). Since the IFDs are the FDs of $s(\theta(t))$, they have the same properties as FDs concerning translation, scaling, rotation and changes of starting point (see sec. 8.3.1.1). optimising $d_{MSD}(s(\theta(.)),\, s'_N(.))$ is equivalent to optimising

$$e_N = d_{MSD}(s(.),\, s'_N(\gamma(.))) \tag{8.3.3}$$

This approach has the advantage that it does not need to interpolate between two contour points $s(t_k)$ and $s(t_{k+1})$. This results in $s^*_N(t) = s'_N \circ \gamma^*(t)$, i.e.

$$s^*_N(t) = \sum_{n=-N}^{N} c_n\, e^{jn\gamma^*(t)} \tag{8.3.4}$$

with $\gamma^*(t)$ the optimal warping function.

Figure 8.5 shows the effect of this warping optimisation. The left column shows the original shapes, the middle column shows the shape reconstruction using 19 FDs, whereas the right column shows the reconstruction result using 19 IFDs. Note that where the left column requires 512 contour points to represent a shape, the second and right columns are reconstructed using only 19 complex numbers. IFDs are generally more effective in the presence of crevices or protrusions (Fig 8.5.C and L) and sharp angles (Fig. 8.5.I and L). Overall the IFD reconstruction shows the leaf incisions more clearly.

8.3.2.1 Optimisation

In order to minimise eq. (8.3.3), an iterative line search algorithm can be used [Philips 95, Boyd 04]. The warping function, $\gamma(.)$, is updated in the n^{th} iteration using the following update rule:

$$\gamma^{(n)}(t) = \gamma^{(n-1)}(t) + \alpha^{(n)} d^{(n)}(t)$$

Here $d(t)$ is the difference between the average error and the error at point t. This results in optimising $\gamma(.)$ by minimising the average points-wise error. The parameter $\alpha^{(n)}$ is the step size that minimises e_N (see eq. (8.3.3)). As an initial value for $\gamma(.)$ we choose $\gamma^{(0)}(t) = t$. There is no proof the algorithm converges in a finite number of iterations. In practice however, e_N generally no longer changes significantly after three iterations.

This iterative process requires multiple calculations of $s_N(\gamma(t))$ and the error, e_N. This can be simplified and speeded up by decomposing $e^{jn\gamma^*(.)}$ in orthonormal polynomials [Philips 96]. The following set of orthonormal polynomials can be used:

$$\begin{aligned} \phi_0(t;\gamma(.)) &= 1 \\ \phi_k(t;\gamma(.)) &= \psi_k(t;\gamma(.)) - \sum_{i=1}^{4} \varphi_{i,k}(t,\gamma(.)) \frac{\phi_{k-i}(t;\gamma(.))}{\|\phi_{k-i}(.;\gamma(.))\|} \end{aligned} \quad (8.3.5)$$

with

$$\psi_k(t;\gamma(.)) = \begin{cases} \cos(\gamma(t))\,\phi_0(t;\gamma(.)) & \text{if } k = 1 \\ \sin(\gamma(t))\,\phi_0(t;\gamma(.)) & \text{if } k = 2 \\ \cos(\gamma(t))\,\phi_{k-2}(t;\gamma(.)) & \text{otherwise} \end{cases} \quad (8.3.6)$$

and

$$\varphi_{i,k}(t;\gamma(.)) = \begin{cases} 0 & \text{if } i > k \\ \langle \psi_k(.;\gamma(.)),\, \phi_{k-i}(.;\gamma(.)) \rangle & \text{otherwise} \end{cases} \quad (8.3.7)$$

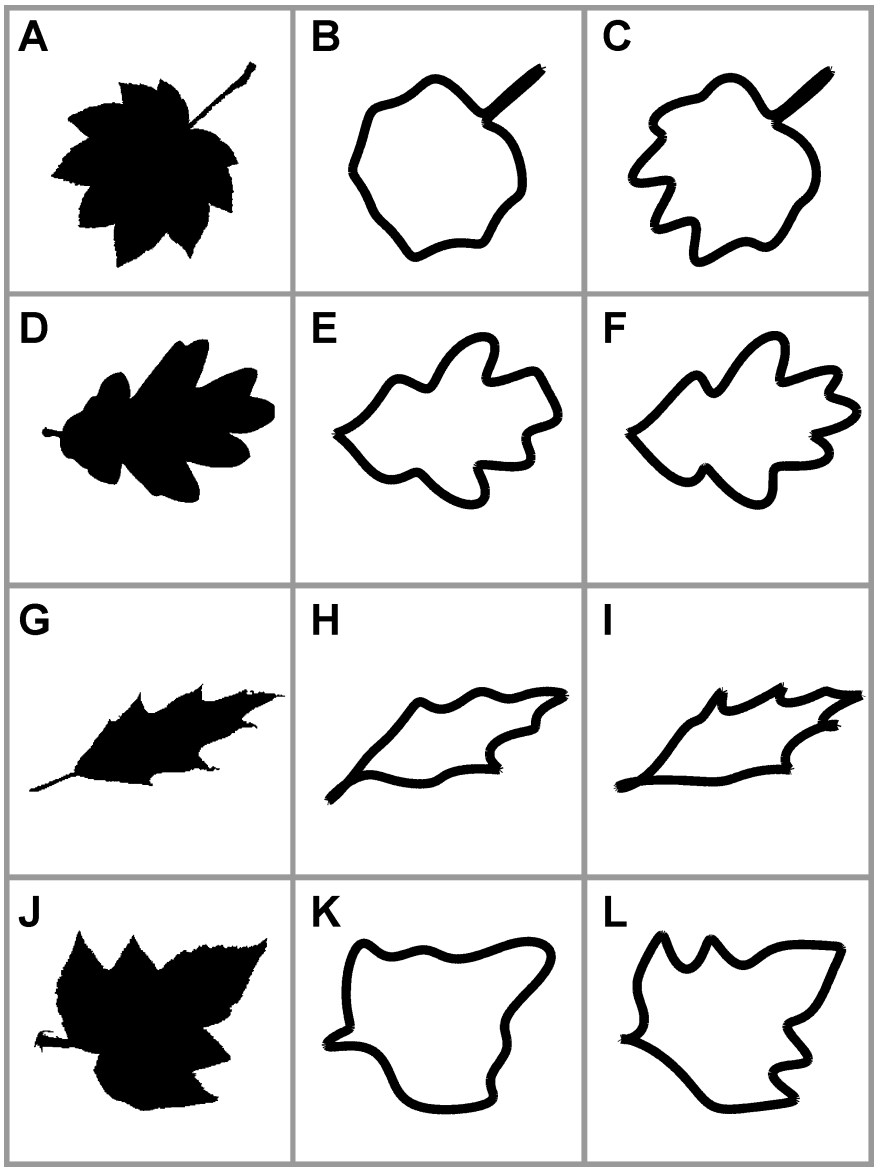


Figure 8.5: Examples of the representation using FDs and IFDs. In the first column, the original shapes are shown. In the second column the approximation using 19 FDs are shown, the last column shows the approximation result using 19 IFDs.

This procedure for orthonormal polynomials has some desirable properties over other orthonormalisation methods, e.g. Gram–Schmidt:

- These orthonormal polynomials are real-valued, and can be computed without complex arithmetic.
- They can be generated efficiently using a four-term recurrence relation.
- The calculation of the bases is independent of the number of bases needed. This is an advantage because the number of polynomials can be increased without having to recompute all the bases.

Note that while the optimisation itself can be performed using orthonormal basis functions, the actual IFD basis functions are not orthonormal.

8.3.3 Shape description using a frame

The previous sections focussed on compressing relevant shape information by using a small number of basis vectors. These basis functions are designed such that no redundancy is encoded due to the orthogonality of the basis vectors. However literature from the last decades has proven that many interesting signal processing results can be achieved by not removing this redundancy [Pižurica 02, Rooms 05b, Goossens 10, Goossens 11, Aelterman 13], instead favouring sets of linear dependent vectors, i.e. frames. These successes, specially using multi-scale frames, triggered the development of a number of different shape descriptors based on frames. Some of the most popular approaches comprise curvature scale space methods [Abbasi 97, Abbasi 99, Gautama 03, Mokhtarian 11], and wavelet based methods [Trimeche 00, Yadav 07]. To overcome the large dimensionality of these frame representations this frame representation is only used as an intermediate result. Two main approaches exist to process this frame representation. The first group does not use the full frame, but makes a selection of vectors of the frame. This frame vector subset depends on the shape itself, where the selected frame vectors are part of the shape descriptor [Mokhtarian 97, Kindratenko 03]. The second approach calculates a set of statistics, e.g. fractional dimension, over the coordinates of the shape relative to the frame [Backes 12, Florindo 12].

8.4 Experimental results

To quantitatively validate different aspects of shape descriptors, a database with images of tree leaves was used. This database contains pictures of isolated leaves

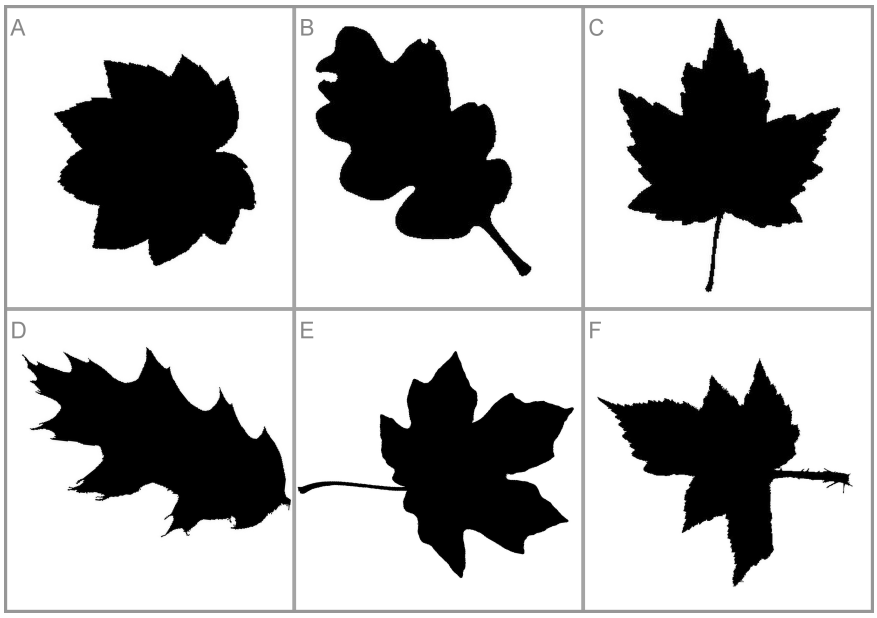


Figure 8.6: Examples of images from the leaf database. This database contains leaves from different trees: A) Circinatum, B) Garryana, C) Glabrum, D) Kelloggii, E) Macrophyllum and F) Negundo.

from six different plant species (*Circinatum, Garryana, Glabrum, Kelloggii, Macrophyllum* and *Negundo*). These six species belong to the genus Acer, i.e. maples (Circinatum, Glabrum, Macrophyllum and Negundo), or to the genus Quercus, i.e. Oak (Garryana and Kelloggii). By selecting species from the same genus, we have a data set with low interclass variability. This results in a realistic but challenging test set for pattern recognition purposes. Fig. 8.6 shows an example for each of these leaf types. In total 436 of these shapes were used to test the different shape descriptors. More information on the dataset can be found in appendix A.

8.4.1 Shape representation

In a first experiment we validate the reconstruction accuracy using a limited number of descriptors. This is relevant for applications where shape descriptors are expected to generate a shape, such applications comprise model based segmentation and shape interpolation. It is also interesting for phenotyping, since it allows a user to visualise a phenotype with a specific set of shape descriptors.

Table 8.1 shows the average Hausdorff distance (see appendix B) between the original shape and the reconstructed shape using different shape descriptors and

using a different amount of shape descriptors. In contrast to what might be expected the Hausdorff distance does not decrease monotonically when more Zernike descriptors are used. This has two main causes, first the Hausdorff distance is a measurement based on the supremum of the minimal distance between sets of contour points, which is sensitive to outliers, if one pixel is wrongfully classified as part of the shape, it can alter this measurement drastically. A second reason is the relatively low number of descriptors used. In Fig. 8.3.E-F the reconstruction of the leaf resulted in little more than an ellipsoid, while that reconstruction also used moments up to the order of 13, which is the highest number of moments used in this experiment. More accurate reconstruction required many more Zernike descriptors. This is not only the case for the example from Fig. 8.3, but generally holds.

The FDs and IFDs are less hampered by outliers, which can be seen by the monotonic decreasing of the Hausdorff distances when using more shape descriptors. The performance of the FDs depends on the kind of shapes they have to represent, e.g., the class of Kelloggii leaves are represented much worse than the class of Garryana leaves. The IFDs perform better than the FD reconstructions for almost all cases, i.e. for all combinations of plant species and number of used shape descriptors. For a low number of shape descriptors this improvement is very small, but for higher amounts of shape descriptors this results in significant improvements, i.e. an average Hausdorff distance of only 39% of the Hausdorff distance resulting from reconstruction with FDs.

Table 8.2 shows the average modified Hausdorff distance, which in contrast to the original Hausdorff distance, is based on an average measurement instead of a supremum (see appendix B for more details). Note that the use of the modified Hausdorff distance doesn't replace the validation based on the original Hausdorff distance. Instead it gives different insight in the performance of shape representation: the modified Hausdorff distance represent how well a contour on average resembles the original shape, whereas the original Hausdorff distance gives an estimate of what the maximal deviation of the shape is. For this measurement the Zernike moments show a better result if moments of higher order are used. However, based on both original and on the modified Hausdorff distance, the Zernike moments perform worse than Fourier contour based descriptors. In terms of modified Hausdorff distance the IFD based shape reconstruction outperforms both Zernike moments and FD based shape reconstruction.

Fig. 8.7 shows the histograms of the relative improvement in modified Hausdorff distance using IFDs to reconstruct the shape instead of the original FDs. The last histogram incorporates all measurements, i.e., for 11, 19, ..., 55 IFDs. A small number of contours are better approximated using the FDs instead of using the IFDs. This is due to the optimisation criteria used for optimising the warping func-

shape descriptors	#descriptors	Circinatum	Garryana	Glabrum	Kelloggii	Macrophyllum	Negundo	Average
ZM	12 (order 5)	63.21	86.75	78.85	114.14	86.26	75.18	86.82
	20 (order 7)	63.48	55.81	85.13	55.99	86.21	66.15	68.46
	30 (order 9)	71.69	59.72	93.95	66.89	97.79	61.20	76.03
	42 (order 11)	66.75	72.18	90.97	73.16	91.06	63.43	77.49
	56 (order 13)	63.07	58.93	84.13	71.84	81.45	52.04	70.22
FD	11	33.38	30.46	25.98	**41.66**	**36.39**	**26.00**	33.30
	19	19.54	15.52	16.57	23.56	20.85	14.83	18.99
	29	11.90	11.23	12.23	16.26	14.18	11.97	13.21
	41	11.25	9.95	11.51	13.29	12.22	12.52	11.78
	55	11.20	9.92	12.26	12.52	11.59	13.45	11.69
IFD	11	**30.30**	**29.19**	**24.59**	41.71	36.79	26.37	**32.48**
	19	**17.79**	**13.85**	**15.00**	**22.81**	**19.63**	**13.35**	**17.62**
	29	**9.64**	**7.35**	**9.62**	**14.26**	**11.48**	**7.74**	**10.39**
	41	**6.68**	**4.28**	**6.68**	**9.19**	**6.82**	**5.05**	**6.66**
	55	**4.70**	**2.67**	**4.84**	**6.22**	**4.44**	**3.80**	**4.55**

Table 8.1: The Hausdorff distance between the original and the reconstructed shape using a different number of shape descriptors. The closer to zero the better (ZM = Zernike moments, FD = Fourier descriptors, IFD = imroved Fourier descriptors).

shape descriptors	#descriptors	Circinatum	Garryana	Glabrum	Kelloggii	Macrophyllum	Negundo	Average
ZM	12 (order 5)	14.02	27.81	22.88	41.54	25.03	20.72	26.80
	20 (order 7)	13.55	14.94	20.39	18.77	22.84	15.39	18.00
	30 (order 9)	13.20	14.10	22.67	18.76	25.54	12.70	18.44
	42 (order 11)	11.32	12.15	19.46	16.01	20.16	10.41	15.45
	56 (order 13)	9.67	9.23	15.06	13.60	15.30	8.39	12.30
FD	11	10.49	10.01	7.52	13.95	11.75	7.56	10.64
	19	5.26	3.99	3.98	5.81	5.41	3.63	4.81
	29	2.45	2.14	2.58	3.42	3.08	2.15	2.72
	41	1.77	1.40	1.89	2.16	1.98	1.57	1.83
	55	1.31	1.02	1.56	1.56	1.53	1.24	1.38
IFD	11	**9.16**	**9.13**	**7.14**	**13.49**	**10.94**	**7.17**	**9.92**
	19	**4.53**	**3.55**	**3.71**	**5.48**	**5.15**	**3.30**	**4.42**
	29	**2.32**	**1.76**	**2.35**	**3.13**	**2.80**	**1.71**	**2.43**
	41	**1.50**	**0.98**	**1.60**	**1.87**	**1.62**	**1.14**	**1.49**
	55	**1.04**	**0.66**	**1.13**	**1.26**	**1.11**	**0.88**	**1.03**

Table 8.2: The modified Hausdorff distance between the original and the reconstructed shape using a different number of shape descriptors. The closer to zero the better (ZM = Zernike moments, FD = Fourier descriptors, IFD = imroved Fourier descriptors).

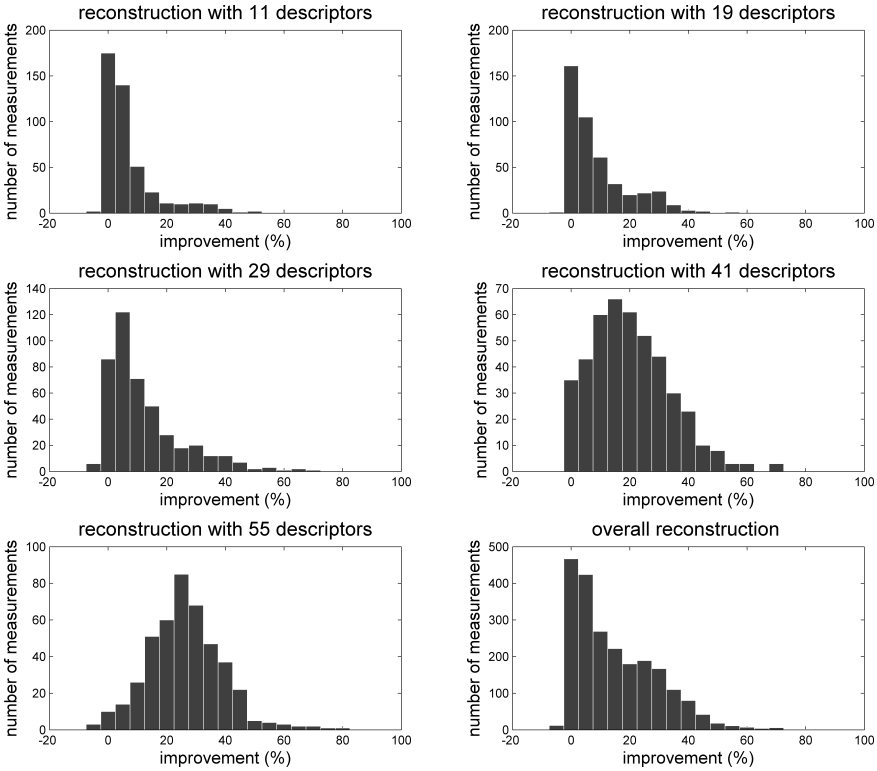


Figure 8.7: Histograms of the improvement of the errors of the IFD approximation over the FD approximation, measured with the modified Hausdorff distance.

tion, which is different from the modified Hausdorff distance. Hence, the warping function is not guaranteed to result in a reconstruction with a lower modified Hausdorff distance. However for the majority of the reconstructed shapes in our database, this method does result in a lower modified Hausdorff distance. We also observe that shape reconstruction using a small set of shape descriptors has less benefit from the non-uniform sampling compared to shape reconstruction using more shape descriptors. However, irrespective of the number of shape descriptors used, the improved Fourier descriptors generally results in better reconstruction. An average improvement of 15% in modified Hausdorff distance is achieved.

8.4.2 Shape classification

In a second experiment we validate the shape descriptors based on their ability to classify shapes. In this experiment we use a maximum likelihood classifier:

$$\hat{c} = \arg\max_{c} \mathrm{p}(\mathbf{s} \mid c), \tag{8.4.1}$$

where $\mathbf{s}$ represents the shape descriptors of the shape that has to be classified, and c iterates over all possible classes, in this situation the possible plant species. The likelihood of a shape descriptor for a specific class is modelled using a multivariate normal distribution. To avoid over-fitting, the parameters of this distribution are trained using a leave-one-out strategy, i.e. for each classified shape the parameters of these distributions are retrained using all shapes except the shape that has to be classified. The set of Negundo images was too small to reliably estimate these parameters, and are therefore omitted in this experiment.

Table 8.3 shows the confusion matrix for the classification using different shape descriptors. This table shows the classification result using 19 shape descriptors that are invariant to rigid transforms, i.e. up to the 7th order in the case of Zernike moments. Note that the Fourier descriptors outperform the Zernike moments. The improved Fourier descriptors generally perform better than the classical Fourier descriptors. The advantage of IFDs results in a better distinction of oak leaves, i.e. Garryana and Kelloggii.

Table 8.4 shows the overall classification results. Given these results it is clear that IFDs result in more accurate classification, again proving the added value of optimising the warping function.

8.4.3 Image content retrieval

In a third experiment we test the applicability of shape descriptors for image content retrieval (ICR) applications. In ICR systems the goal is to retrieve content, in our case shapes, similar to a specific input shape. For this application dissimilarity is expressed using the Euclidean distance between the shape descriptors from the input shape and the shape descriptors in the database. So the most similar shapes will have the lowest Euclidean distance. Fig. 8.8 shows an example of the five shapes that are the most similar to the shape in Fig. 8.6.A. This input shape corresponds to a Circinatum leaf, therefore we chose to show all retrieved shapes that also correspond to a Circinatum leaf in black, whereas leaves corresponding to other trees are illustrated in grey. The first row corresponds to the shapes with the lowest Euclidean distance between its shape descriptors and the shape descriptors from the input shape, i.e. so the most similar shapes, the second row with the

			predicted class				
			Circinatum	Garryana	Glabrum	Kelloggii	Macrophyllum
ZM	real class	Circinatum	59	0	0	0	2
		Garryana	1	12	1	12	1
		Glabrum	0	0	43	2	14
		Kelloggii	0	2	0	86	1
		Macrophyllum	2	0	11	1	39
FD	real class	Circinatum	59	0	0	1	1
		Garryana	0	13	0	10	4
		Glabrum	0	0	**57**	0	2
		Kelloggii	0	0	0	**89**	0
		Macrophyllum	0	1	1	0	51
IFD	real class	Circinatum	**60**	1	0	0	0
		Garryana	0	**16**	0	9	2
		Glabrum	0	0	56	1	2
		Kelloggii	0	0	0	**89**	0
		Macrophyllum	0	0	0	0	**53**

Table 8.3: The confusion matrix for shape classification using the following shape descriptors: Zernike moments (ZM), Fourier descriptors (FD) and improved Fourier descriptors (IFD).

	Precision			Recall			F-measure		
	ZM	FD	IFD	ZM	FD	IFD	ZM	FD	IFD
Circ.	0.95	1	**1**	0.97	0.97	**0.98**	0.96	0.98	**0.99**
Garr.	0.86	0.93	**0.94**	0.44	0.48	**0.59**	0.59	0.63	**0.83**
Glab.	0.78	0.98	**1**	0.73	**0.97**	0.95	0.75	**0.97**	**0.97**
Kell.	0.85	0.89	**0.90**	0.97	1	**1**	0.90	0.94	**0.95**
Macr.	0.68	0.88	**0.93**	0.74	0.96	**1**	0.71	0.92	**0.96**

Table 8.4: Classification results using different shape descriptors. The closer these values are to one, the better the classification (ZM = Zernike moments, FD = Fourier descriptors, IFD = imroved Fourier descriptors).

second lowest distance, etc. Both Zernike moments and Fourier based descriptors result in a shape corresponding to a Circinatum leaf as the best match, i.e. at the first position. However the IFDs are the only shape descriptors resulting in only Circinatum leaf shapes in the top five matches. The regular FDs have two Glabrum leaves in the top five. These leaves show some similarity to Circinatum leaves, but of course it would be better to only have Circinatum leaves as best matches. The retrieval using Zernike moments has one Glabrum leaf in it's top five (the third best match), but also has two Macrophyllum leaves in the top five, which show strong recesses that are not present in the input shape.

The performance of an ICR system can be quantified through the true positive rate (TPR) and false positive rate (FPR), respectively the number of relevant shapes retrieved over the number of shapes retrieved and the number of irrelevant shapes (shapes from different trees than the input query) over the number of retrieved shapes. By changing the number of retrieved shapes, a trade-off can be made between both measurements. The Receiver Operating Characteristics (ROC) is a graphical plot of TPR as a function of FPR. Figure 8.9 shows the average ROC curves corresponding to shape retrieval using ZMs, FDs and IFDs. These ROC curves show the performance of shape retrieval using 65 descriptors, this is equivalent to ZMs up to order 14. The higher the area under the ROC curve (AUC), the better the performance. Zernike moments perform significantly worse than FDs, i.e. an AUC of 0.62 compared to an AUC of 0.83. IFDs perform best with an AUC of 0.85. Curvature scale space (CSS), as part of the MPEG standard, has been pitched as the tool for image content retrieval [Mokhtarian 97, Abbasi 99, Mokhtarian 11], however literature already reports several experiments showing that both Zernike moments and Fourier descriptors are better suited for this task [Zhang 03, El-ghazal 09, Amanatiadis 11]. Therefore, we have omitted the comparison with CSS in this experiment.

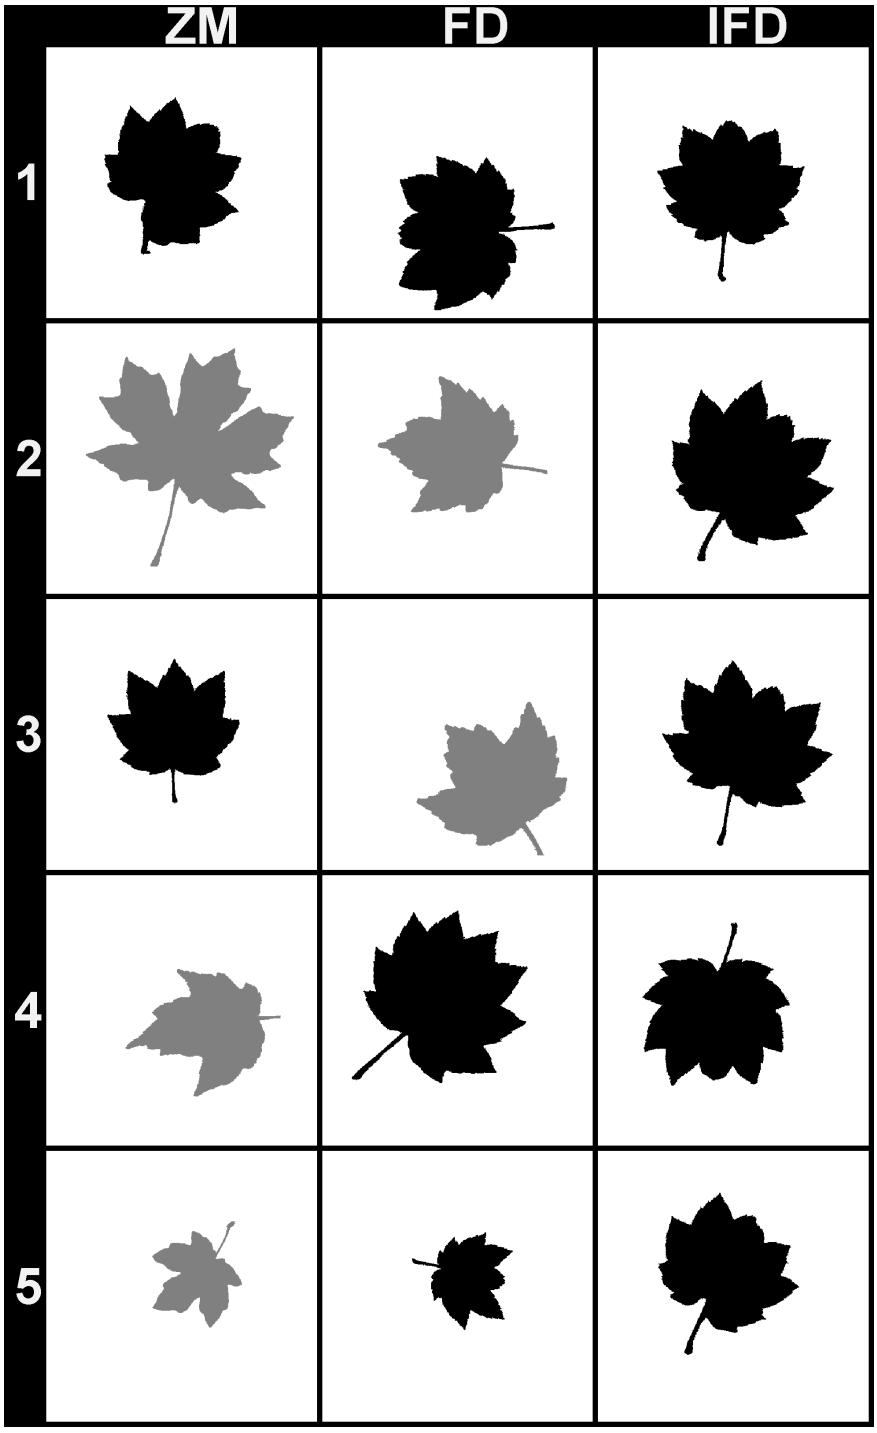


Figure 8.8: The five best matches for the Circinatum query from Fig. 8.6.A. The rows correspond to the n^{th} best match, whereas the columns correspond to the different shape descriptors used, i.e. respectively Zernike moments, Fourier descriptors and improved Fourier descriptors.

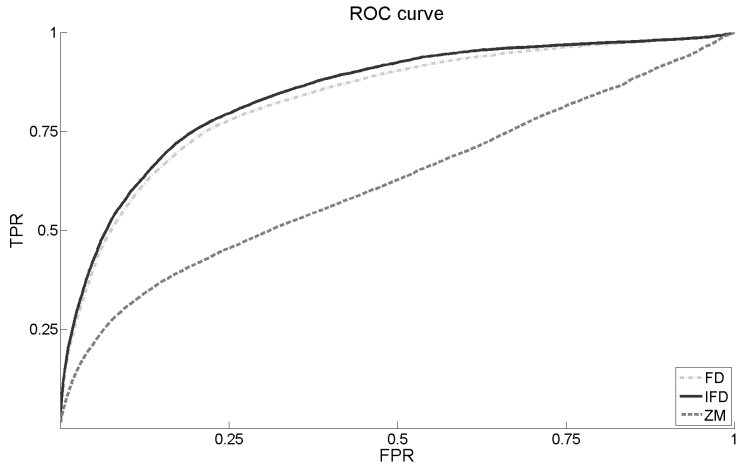


Figure 8.9: Three Receiver Operating Characteristics (ROC) corresponding to correct shape retrievals using different shape descriptors. The higher the area under the ROC (AUC), the better the performance.

8.5 Conclusion

This chapter discussed the problem of quantifying shape. It describes the three major classes of shape descriptors: scalar morphometrics, region based descriptors and contour based descriptors. The chapter gave an overview of state-of-the-art methods for all three classes. Furthermore, we proposed a new contour based shape descriptor: the improved Fourier descriptor. In contrast to other contour based shape descriptors these descriptors are not the result of minimising distance between the reconstruction and a given shape signature. Instead the distance between the actual shape and the reconstruction is minimised. This is achieved using a non-uniform Fourier transform, thus avoiding extensive interpolation of contour points. This new method was compared with state-of-the-art methods for a number of pattern recognition applications: shape reconstruction, shape classification and image content retrieval. For all applications the proposed shape descriptors outperformed the classical Fourier descriptors, and resulted overall as the best method. The improved Fourier descriptors resulted in an average improvement of 15% in modified Hausdorff distance, an average F-measure of 0.94 for classification, compared to 0.89 for state-of-the-art methods and an AUC 0.85 for ICR compared to 0.83 for the classical Fourier descriptors. These results show that the method has generic value for numerous shape description applications.

This work is published in a number of international conference proceedings [De Vylder 08b, De Vylder 08a, De Vylder 09c, De Vylder 12d, Janssens 13], in a

peer reviewed journal publication [De Vylder 12c] and in a journal publication in preparation [De Vylder 14b].

9 Conclusion

The main part of this dissertation addresses the automatic analysis of images or of objects observed in images. Techniques working for both still images, and for the analysis of time lapse sequences have been discussed. Experiments on both simulated and real data show the accuracy and computational efficiency of the different proposed techniques and its practical use in a wide range of applications.

In this chapter, we review our main contributions for image analysis and shape description. This thesis is finally concluded with some hints for future research.

9.1 Review of our contributions

To summarise, the novelties and contributions of this dissertation can be divided in two groups: (1) those related to the automatic detection and delineation of objects of interest in images and (2) those related to describing shape and deformation. In the first group we have several contributions, working for different applications and using different algorithms:

Multiple improvements were made on the active contour framework, which is widely used for segmentation and tracking. Active contours require the minimisation of an energy function, which is typically a linear combination of a data-fit term and regularisation terms. This energy function can be tailored to the intrinsic

object and image features. This is normally done by either modifying the actual terms or by changing the weighting parameters of the terms. There is however no protocol on how to set these terms and optimal weighting parameters for a given application. Although heuristic techniques exist for parameter estimation, often trial and error is used. In Chapter 5, we propose a probabilistic interpretation for active contour segmentation. The proposed Bayesian framework results in a generalisation of state-of-the-art active contour segmentation, e.g. segments can be modelled as smooth, non-constant objects, or containing texture, or segments can have a specific shape, etc. In the proposed framework, all parameters have a statistical interpretation, thus avoiding ad hoc (non-optimal) parameter settings.

Cell nuclei detection in fluorescent microscopic images is an important and time consuming task in a wide range of biological applications. Blur, clutter, bleed through, clustering of nuclei, and partial occlusion of nuclei make individual nuclei detection a challenging task for automated image analysis. In Chapter 6 we propose a number of new segmentation and detection methods for this task. These methods focus on two important problems: non-uniform lighting and detecting individual nuclei in dense clusters of nuclei. Improvement over conventional approaches is achieved by exploiting prior knowledge of the nucleus shape and appearance in relation to its surrounding neighbourhood. An example of such new prior knowledge is the use of a shape dictionary in combination with the convex active contour framework. The proposed methods shows accurate detection results for dense clusters of nuclei, both in 2D images and for 3D datasets.

Automated image analysis of plant rosettes is an important non destructive method for studying plant growth. While several computer vision methods have been proposed for this purpose, these algorithms only work in the restricted environment of some specific monitoring system. We developed a new toolbox that is not constraint to one specific monitoring system and can be adapted to different monitoring setups without the need of extensive user input. The proposed toolbox also analyses photosynthesis, leaf temperature and plant growth related parameters through the analysis of VIS, chlorophyll fluorescence and/or near infrared time lapse sequences.

For the tracking and segmentation of individual leafs in time lapse sequences we extended the parametric active contour framework. The most important novelties that we proposed encompass a computational efficient force field, an optimisation scheme to cope with the low temporal resolution of the sequences, a force field based both on edge and intensity information and a force field incorporating prior knowledge on the shape of the plants. The added value of the proposed adjustments is thoroughly evaluated on real time lapse sequences captured using VIS, fluorescence and IR cameras.

In the second group, the group of novelties related to shape, we have proposed the

following contributions:

We developed a new set of shape descriptors, inspired by the popular Fourier descriptors. By first warping the scanning speed of the contour before calculating the Fourier descriptors, the shape approximation is significantly improved. We tested these new shape descriptors for several applications: classification of shapes, retrieval of similar shapes in a database and shape reconstruction, i.e. to what extent a shape can be reconstructed from a small set of shape descriptors. For each of these applications, the proposed improved Fourier descriptors outperform state-of-the-art shape descriptors.

We investigated the use of an objects boundary to register the objects interior. The proposed approach has the advantage that it works irrespective of the presence of stable landmarks. Furthermore does this method guarantee to only register motion due to the overall motion and deformation of the object, without over-fitting to local motion inherent to small particles within the object. This makes it an interesting method for the analysis of genome dynamics in pathological cell nuclei. The nuclei of pathological cells such as cancer cells deform in an apparently random way. These apparent random deformations make the use of algorithms from literature a non-trivial task, while our proposed method is generic yet powerful enough to cope with these deformations.

9.2 Directions for future research

It is not our intention to list all possible adjustments that could be made on the presented algorithms. Instead we would like to highlight four interesting and challenging directions for future research.

In Chapter 5 we proposed a probabilistic segmentation method. Although all parameters have a statistical interpretation, the correct model has to be chosen and a correct estimation of all variables still has to be made. Hence the threshold for implementation is relatively high, since these methods in general need an expert to initiate the segmentation method. Therefore it would be useful to research the use of statistical estimators to check which model should be used and to estimate the necessary parameters. We believe that especially for the correct setting of the shape priors, it would be interesting to strive for a semi-automatic method. Based on a maximum-likelihood estimator, i.e. so without any shape prior, an initial segmentation result could be achieved. Using machine learning methods, the found segments could be clustered, based on their shape. These clusters can then be checked by a human observer, which requires much less effort then manually annotating an initial training set or tweaking parameters on a trial and error basis.

Based on the statistics of the found clusters, the exact parameters for the priors could be estimated.

In chapter 6 we proposed a segmentation technique based on geometric active contours combined with an energy term based on a shape dictionary. Tests showed that this method results in accurate segmentation of mononuclear blood cells. It would be interesting to expand this method to detect other shapes than circles. This could for example be used to detect elliptical nuclei, or based on specific shape descriptors, less regular shapes could be detected. Although the proposed shape dictionaries are generic, it would become computational impractical to adjust this to shapes that require more then model 3 parameters, e.g. location and size. The following two approaches would be interesting to investigate:

1. First, apply an initial segmentation, based on some confidence measure (conform the method proposed in section 6.2.2) only apply the dictionary based method on clustered segments that are too challenging to process with other methods. This already could limit the parameter space significantly, thus allowing the use of a dictionary based approach.

2. Instead of calculating the dictionary explicitly and using this to approximate the segmentation result, it could also be interesting to investigate transforms that map the segmentation result to a dictionary representation. This would be similar to the analysis approach instead of a synthesis approach used in image restoration methods.

In Chapter 8 we proposed a new set of shape descriptors. By optimising the scanning speed of the contour, better shape descriptors are achieved. We investigated this approach to improve the popular Fourier descriptors. This could however also be applied using other basis functions as well. It would be interesting to test to what extent the observed improvements translate to other shape descriptors. For example, a set of dedicated basis functions could be trained for a specific application, then afterwards the scanning speed optimisation approach could be used for the shape description of shapes outside of the training dataset.

In Chapter 2 we discussed several new emerging image modalities. While many computer vision methods can be translated to these new types of images, some new challenges will arise. One challenge that can be expected is the scalability of existing computer vision algorithms. Many of the new image modalities are able to gather very large datasets, e.g. by capturing more spectra, higher spatial resolution, 3D, higher temporal resolution, etc. One approach to handle this big data, is to investigate efficient and preferably parallel implementations or new algorithms. These could enable the high level parallelism of multi-core central processing units (CPU) or the massive parallelism that can be achieved by graphics processing units

(GPU). Although several other challenges can be expected for these new image modalities, they also form a great opportunity, where (existing) image processing algorithms can enable biological breakthroughs.

A

Data sets

In this appendix we briefly summarise all relevant information on the used datasets.

A.1 Fibroblast cells

A.1.1 Cell culture and transfection

ECV-TRF2 were cultured in advanced DMEM (Gibco, Invitrogen, Merelbeke, Belgium) supplemented with 2% FCS and 1% pencillin/streptomycin/glutamin. ECV-TRF2, is a derivative clone from ECV304, a cell line originally described as an immortalized endothelial cell line [Takahashi 90], stably expressing the telomere binding fusion protein mCitrine-hTRF2 [De Vos 09b]. Fibroblast cells from a patient with a nonsense Y259X homozygous mutation in the lamin A/C gene (LMNA−/−) which are completely devoid of functional lamin A/C [Muchir 03], were cultured in advanced DMEM/F12 (Gibco, Invitrogen, Merelbeke, Belgium) supplemented with 10% FCS and 1% penicillin/streptomycin/glutamin.

For transfections and live cell imaging, cells were grown on collagen coated glass bottom dishes (Matek, USA). LMNA−/− cells were transiently transfected with the construct pmCitrine-hTRF2 and Lipofectamin 2000 (Invitrogen, Merelbeke, Belgium) according to manufacturer's instructions. Transiently transfected cells were imaged 24 h after transfection.

A.1.2 Image acquisition

For live cell imaging, glass bottom dishes were filled with prewarmed medium and sealed off with parafilm, after which they were inserted into a microscope incubator and allowed to equilibrate to 37°C for 30 min. Cells were observed under a Nikon A1r confocal microscope mounted on a Nikon Ti body (Nikon Instruments, France). Experiments were carried out with a Plan Apo 40 × /1.3 oil objective. Two-hour time-lapse recordings of LMNA−/− cells were made at 2 min intervals with minimally invasive CLEM technology [Hoebe 07]. 12-bit images were sampled at 288 × 288 × 10 (x, y, z) voxels, with a physical voxel size of 104 nm × 104 nm × 500 nm. For the validation experiments, ECV-TRF2 cells were chemically fixed by addition of cold (4°C) 100% methanol (Sigma, Belgium) in fluxo, i.e. during acquisition at 30 s intervals. All images by courtesy of Winnok De Vos and Cell Systems & Imaging, Ghent university.

A.2 Hybridoma cells

A.2.1 Cell culture and transfection

OKT3 and HSAspec hybridoma cells where stained with Alexa Fluor 555 targeting against HSA. Cells where embedded in Opti-MEM® I Reduced Serum Media (http://www.lifetechnologies.com/be/en/home/life-science/cell-culture/mammalian-cell-culture/classical-media/opti-mem.html)

A.2.2 Image acquisition

Image acquisition was carried out with the AxioVision multichannel fluorescence module and the AxioCam MRm camera (Carl Zeiss). Cell nuclei were visualized using Zeiss filter set no. 49 (G 365 nm, FT 495, BP 445/50). Slides were scanned at 20× magnification using a Carl Zeiss short distance Plan-Apochromat ® objective [Vandewoestyne 09a]. Images were acquired and were stored as RAW tiff-files. All images by courtesy of Cynthia Lambrecht, Mado Vandewoestyne and Laboratory for Pharmaceutical Biotechnology, Ghent University.

A.3 Embryonic cells

A.3.1 Cell culture and transfection

Medaka embryos were fixed in 4% paraformaldehyde and stained for nuclei with Sytox Green. Zebrafish wild-type embryos were dechorionated at the one-cell stage and injected with 20 μM H2B-mCherry mRNA and 20 μM ras-eGFP mRNA. Embryos of the ras-eGFP transgenic line were dechorionated at the one-cell stage. Stained Medaka embryos and injected or transgenic zebrafish embryos were embedded in 1% low-melting-temperature agarose as previously described6. The DSLM imaging chamber was filled with 30% Danieau's buffer (for live zebrafish imaging) or 1$\times$ PBS (for fixed Medaka embryos).

A.3.2 Image acquisition

DSLM imaging was performed using the 488 nm and 568 nm lines of an Argon-Krypton gas laser (Melles Griot) for illumination and RazorEdge 488 RU and 568 RU long-pass filters (Semrock) in the detection system. For volumetric imaging, images were acquired with a 4-μm-thick light sheet (FWHM) in steps of 1.8μm. The exposure time was set to 50 ms per image for all experiments. The specimen position along the vertical axis was manually corrected to compensate for the drift of the agarose cylinder during the first few hours of imaging. Typically, a total drift distance of up to 50 μm occurred over a period of 3 h before the vertically mounted agarose matrix was fully equilibrated.

The DSLM-SI patterns are created by using the following hardware components: an AA.AOTF.nC-400-650nm-PV-TN AOTF (AA Opto-Electronic), a pair of VM-500+ laser scanners operated via an HC/3 scan controller (GSI Lumonics), an S4LFT0061 F-theta lens with 60-mm focal length (Sill Optics), a tube lens with 164.5-mm focal length (Carl Zeiss), an NI-6733 data acquisition board (National Instruments) and a Plan-Apochromat 5$\times$/0.16 W illumination lens (Carl Zeiss). We equipped the DSLM with the following detection lenses: C-Apochromat 10$\times$/0.45 W, Plan-Apochromat 20$\times$/1.0 W and Plan-Apochromat 63×/1.0 W (Carl Zeiss). For further details on the acquisition we refer to [Keller 10]. The dataset is courtesy of Philipp J. Keller and European Molecular Biology Laboratory, Germany.

A.4 Peripheral blood mononuclear cells

A.4.1 Cell culture and transfection

Peripheral blood mononuclear cells (PBMCs) were isolated from a healthy volunteer's EDTA blood sample by density gradient centrifugation on Ficoll-Paque Plus (GE Healthcare, Diegem, Belgium) according to the manufacturer's instructions. 300.000 PBMCs were cytospun on a Poly-L-lysine coated slide (DAKO) as previously described [Vandewoestyne 09b]. The slides were air dried and fixed for 1 minute in 70% EtOH. After air-drying, the slides were mounted with antifade Vectashield mounting solution (Vector Labs, Burlingame, CA, USA) containing 4´,6-diamidino-2-phenylindole dihydrochloride (DAPI, 400 ng/ml, Sigma-Aldrich) to counterstain all nuclei on the slide. A coverslip was applied.

A.4.2 Image acquisition

Image acquisition was carried out with the AxioVision multichannel fluorescence module and the AxioCam MRm camera (Carl Zeiss). Cell nuclei were visualized using Zeiss filter set no. 49 (G 365 nm, FT 495, BP 445/50). Slides were scanned at 20× magnification using a Carl Zeiss short distance Plan-Apochromat ® objective [Vandewoestyne 09a]. Images were acquired and were stored as RAW tiff-files. All images by courtesy of Mado Vandewoestyne, Trees Lepez and Laboratory for Pharmaceutical Biotechnology, Ghent University.

A.5 VIS images A. thaliana

A.5.1 Ecotypes and growth

Arabidopsis Col-0 wild type seeds were acquired from the Nottingham Arabidopsis Stock Center (Nottingham, UK). Plants were grown for 4 weeks in a growth room with 22°C and 16h/8h light/dark cycles. Light intensity was 60µmolesm^{-2}s^{-1} from cool white light tubes (Philips, Eindhoven, The Netherlands).

A.5.2 Image acquisition

A photograph was taken using a Canon 500D reflex camera with 18-55 mm lens (www.Canon.com), at 55mm, from a distance of 50 cm. All images by courtesy of Dominique Van Der Straeten and Laboratory of Functional Plant Biology, Ghent university.

A.6 DIC images A. thaliana

A.6.1 Ecotypes and growth

Arabidopsis Col-0 seeds for time-lapse analysis were sterilized in 3% bleach for 15 min and sown on medium containing 0.5x MS medium (Duchefa) solidified with 0.9 g L-1 plant tissue culture agar (Lab M) on round Petri dishes. After a stratification period of 2 days, the plates were placed in a growth chamber under long-day conditions (16 h of light, 8 h of darkness) at 21°C. Average light intensity, supplied by cool-white fluorescent tubes (Spectralux Plus 36W/840; Radium), was around 80 mE m-2 s-1. Plant were harvested after 20 days. Leaves were cleared in ethanol and lactic acid before mounting on microscopy slides.

A.6.2 Image acquisition

Imaging was performed using the Olympus BX53 microscope equipped with DIC optics. Images were acquired with an Olympus SC30 camera. All images by courtesy of Stijn Dhondt, department of Plant Systems Biology, Ghent university - VIB.

A.7 Chlorophyll fluorescence and IR images A. thaliana

A.7.1 Ecotypes and growth

Arabidopsis accessions and *abi1-1* were acquired from the Nottingham Arabidopsis Stock Center (Nottingham, UK). 35S::ABI5 was kindly provided by J. Smalle. Plants were grown in trays for 3 weeks in 13h/11h light/dark cycles in a growth room at 21°C. Daytime PPFD was 150 $\mu molm^{-2}s^{-1}$.

A.7.2 Image acquisition

Rosettes were photographed every hour by a robotized camera system [Chaerle 06]. Fluorescence images were acquired with an in-house developed fluorescence imaging system; [Chaerle 04], while infrared images where acquired with a FLIR-AGEMA Thermovision THV900SWTE (Flir systems, Croissy-Beaubourg, France). All images by courtesy of Dominique Van Der Straeten and Laboratory of Functional Plant Biology, Ghent university.

plant	Circinatum	Garyana	Glabrum	Kelloggii	Macrophyllum	Negundo	complete dataset
# images	64	84	75	94	82	37	436

Table A.1: An overview of the number of images for each type of leaf.

A.8 Chlorophyll fluorescence and IR images of sugar beet seedlings

A.8.1 Ecotypes and growth

All plant material is property of Erik De Bruyne, SESVanderhave NV.

A.8.2 Image acquisition

The seedlings were photographed every hour by a robotized camera system [Chaerle 06]. Fluorescence images were acquired with an in-house developed fluorescence imaging system; [Chaerle 04], while thermal infrared images where acquired with a FLIR-AGEMA Thermovision THV900SWTE (Flir systems, Croissy-Beaubourg, France). All images by courtesy of Dominique Van Der Straeten and Laboratory of Functional Plant Biology, Ghent university.

A.9 Digital herbarium

A.9.1 Ecotypes and growth

This database contains pictures of isolated leafs from six different plant species (*Circinatum, Garryana, Glabrum, Kelloggii, Macrophyllum and Negundo*). In total 436 of these images were used to test the use of IFD's for shape approximation. The database is courtesy of T.G. Dietterich and the Plant and Botany Department, Oregon State University, for providing the digital herbarium.

A.9.2 Image acquisition

These are colour pictures of 512×512 pixels. Each image contains an isolated leaf on a white background. The leafs were extracted out of the picture using a

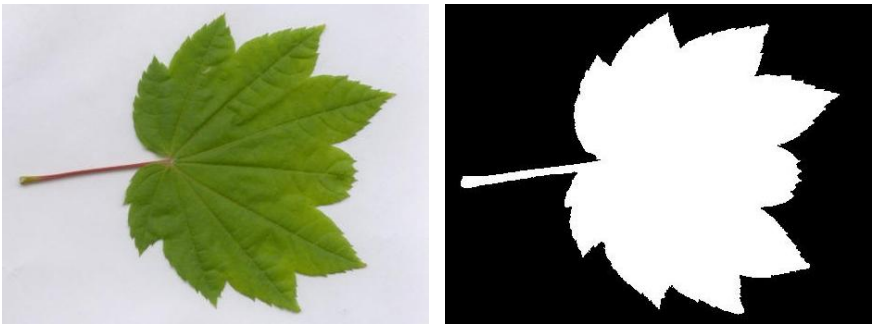


Figure A.1: An example of segmentation of leafs using a simple threshold on the RGB values.

threshold on the RGB values: all pixels with value below 200 for the red channel where classified as leaf pixels. Fig. A.1 shows an example of this colour filter.

B

Quality metrics

B.1 Detection metrics

In order to validate the methods that have to detect certain objects, e.g. objects of interest, we use a number of validation metrics:

- Root mean squared error (RMSE) is an error measure on the count of objects:
  $RMSE = \sqrt{\frac{1}{M}\sum_{i=1}^{M}(D(i) - C(i))^2}$, with $D(.)$ the number of detected objects in the i^{th} image of the dataset, whereas $C(.)$ represents the real number of obecects of interest in the i^{th} image of the dataset.
- True positives (TP): the number of objects of interest in the ground truth that are correctly detected using the proposed algorithm.
- False positives (FP): the number of detections which have no corresponding ground truth object of interest.
- False negatives (FN): the number of objects of interest in the ground truth with no corresponding detection.
- Precision (P) is the ratio of the number of correct detected objects of interest (TP), over the total number of detected objects of interest, i.e. including false positives.

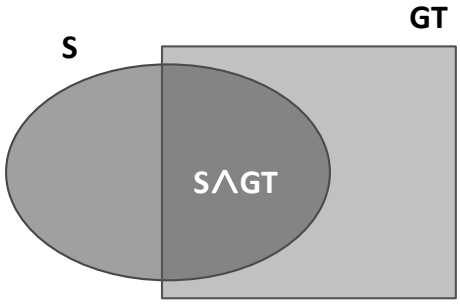


Figure B.1: An example of a potential segmentation output, with S the found segment, and GT corresponding to the ground truth of the same segment.

- Recall (R) is the ratio of the number of TP's in an image, over the total number of objects of interest in the ground truth.
- the F-measure takes both false positives and false negatives into account by combining precision and recall: $F = 2\frac{\text{precision}\cdot\text{recall}}{\text{precision}+\text{recall}}$.

B.2 Segmentation metrics

The Dice coefficient is often used as a metric for the validation of segmentation results [Dice 45]. Fig. B.1 shows an example of a potential segmentation output, with S is the result of a segmentation algorithm,e.g. the region enclosed by a parametric active contour, and GT the ground truth segment. The Dice coefficient between S and GT is defined as:

$$d(\mathrm{S},\mathrm{GT}) = \frac{2\ \mathrm{Area}(\mathrm{S}\wedge\mathrm{GT})}{\mathrm{Area}(\mathrm{S})+\mathrm{Area}(\mathrm{GT})}, \tag{B.2.1}$$

where $S \wedge GT$ consist of all pixels which both belong to the detected segment as well as to the ground truth segment. If S and GT are equal, the Dice coefficient is equal to one. The Dice coefficient will approach zero if the regions hardly overlap.

B.3 Tracking metrics

For the validation of tracking methods we use the false matches (FM) metric. This counts the number of detected objects that are incorrectly matched with detections in the previous time frame. Which matching algorithm is used depends on the application and or proposed technique. The actual used matching algorithm will be discussed in the section explaining each technique or application.

B.4 Shape difference metrics

In order to compare different shape representations, several metrics to express similarity or dissimilarity between shapes have been proposed. The mean squared distance is a metric which is used to compare contour parametrization. Consider s_1 and s_2 to be two parametrized curves describing two shapes, the mean squared distance between s_1 and s_2 is defined as follows:

$$d_{msd}(s_1, s_2) = \frac{1}{N} \sum_{k=0}^{N-1} (s_1(t_k) - s_2(t_k))^2. \tag{B.4.1}$$

This might be a good measurement for differences between functions, but it is not necessarily a good measurement for difference between shapes. Let s_1 and s_2 describe the same shape, but with an other starting point, e.g. $s_2(t_k) = s_1(k+1)$ with $t_N = t_0$. Even though the shapes are equal, the mean squared distance differs from zero. To overcome this problem, two other metrics are defined in [Zhang 05]: the Hausdorff distance and the modified Hausdorff distance. For two sets of contour points $A = a_1, a_2, ..., a_p$ and $B = b_1, b_2, ..., b_q$ the Hausdorff distance (HD) is defined as follows

$$d_H(A, B) = \max(h(A, B), h(B, A)), \tag{B.4.2}$$

where

$$h(A, B) = \max_{a \in A} \min_{b \in B} d(a, b).$$

With d a given metric function. For our purpose the euclidean distance is used. Fig. 1 shows an example where it is clear that $h(A, B)$ is not equal to $h(B, A)$. If we would like to have the Hausdorff distance between the ellipse and the contour of the heart, we need h(ellipse,hart) and h(hart,ellipse). The length of the dashed arrow is equal to h(ellipse,hart), where h(hart,ellipse) equals the length of the dotted arrow.

The Hausdorff distance is a measurement of the maximum distance between two contours. Whereas the modified Hausdorff distance (MHD) is a measurement of the average distance between two contours. This distance function is given by eq. (B.4.2) where

$$h(A, B) = \frac{1}{p} \sum_{a \in A} \min_{b \in B} d(a, b).$$

The Euclidean distance is used for d.

B.5 Classification metrics

For the classification of shapes we use the following quality metrics:

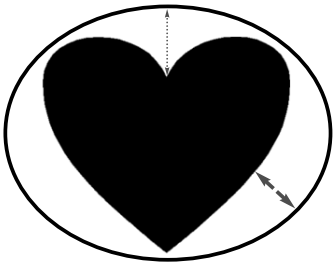


Figure B.2: An example of the non-symmetrical nature of h(.,.) in the Hausdorff distance.

- Precision (P) is the ratio of the number of correct classifications (TP) for a specific class, over the total number of object classified as that class, i.e. including false positives.
- Recall (R) is the ratio of the number of correct classification for a specific class, over the total number of objects of the predefined class in the ground truth.
- the F-measure takes both false positives and false negatives into account by combining precision and recall: $F = 2\frac{\text{precision}\cdot\text{recall}}{\text{precision}+\text{recall}}$.

C

Curriculum Vitae

Personalia

Name Jonas De Vylder

Adress: Jutestraat 7w304, 9160 Lokeren, Belgium

Email: jonas.devylder@telin.ugent.be

Mobile: +32 498 77 82 34

Place of birht: Sint niklaas, Belgium

Nationality: Belgian

Education

2007-2014 Ph.D. fellowship, Department of telecommunications and information processing, Ghent university, Ghent, Belgium.

2003-2007: Master in Informatics, option formal informatics; obtained at Ghent university, Ghent, Belgium.

2006-2007: Undergraduate dissertation, "Characterization and recognition of cell structures in biomedical images", promoters: prof. dr. ir. Wilfried Philips, prof. dr. ir. Sidharta Gautama.

2008: Advanced Statistics and Data Mining Summerschool, Universidad CEU San Pablo, Madrid, Spain.

Awarded grants

Ph.D. fellowship of the agency for innovation through science and technology (IWT) January 2008- December 2011; poject title: "Segmentatie en morfologische analyse van celkernen"

publications

1. Goossens B, De Vylder J, Philips W. Quasar - a new Heterogeneous Programming Framework for Image and Video Processing Algorithms on CPU and GPU. Proc. of the International Conference on Image Processing (ICIP2014). Paris, France; 2014

2. Xie X, De Vylder J, Van Cauwelaert D, Veelaert P; Philips W, Aghajan H. Average Track Estimation of Moving Objects Using RANSAC and DTW. ACM/IEEE International Conference on Distributed Smart Cameras (ICDSC). Venice Italy; 2014

3. De Vylder J, Rogiest W. Extra winst voor grote partijen: wiskunde van verkiezingen. EOS. 2014;5:34.

4. De Vylder J, Van Haerenborgh D, Aelterman J, Philips W. A probabilistic interpretation of Geometric Active Contour Segmentation. European Signal Processing Conference, Proceedings. Lisbon, Portugal; 2014.

5. De Vylder J, Luong Q, Dhondt S, Vlaminck M, Inzé D, Philips W. Rotation and focal depth fusion of DIC microscopic images. IEEE International Symposium on Biomedical Imaging, Abstracts. 2014.

6. Roels J, Aelterman J, Luong Q, De Vylder J, Philips W, Lippens S, et al. Image Optimization for 3D Electron Microscopy Images. 2014.

7. Aelterman J, De Vylder J, Luong Q, Lippens S, Saeys Y, Goossens B, et al. The image processing revolution: the next step in electron microscopy. 3View Users Group Meeting and Symposium, Abstracts. 2014.

8. Babin D, Pizurica A, De Vylder J, Vansteenkiste E, Philips W. Brain blood vessel segmentation using line-shaped profiles. PHYSICS IN MEDICINE AND BIOLOGY. 2013;58(22):8041-61.

9. De Vylder J, Van Der Straeten D, Philips W. Multiple leaf tracking using computer vision methods with shape constraints. Proceedings of SPIE, the International Society for Optical Engineering. Bellingham, WA, USA: SPIE, the International Society for Optical Engineering; 2013.

10. De Vylder J, Rooms F, Dhondt S, Inzé D, Philips W. A novel tracing method for the segmentation of cell wall networks. IEEE Engineering in Medicine and Biology Society Conference Proceedings. New York, NY, USA: IEEE; 2013. p. 5433-6.

11. Aelterman J, Goossens B, De Vylder J, Pizurica A, Philips W. Computationally efficient locally adaptive demosaicing of color filter array images using the dual-tree complex wavelet packet transform. PLOS ONE. 2013;8(5):1-18.

12. De Vylder J, Aelterman J, Lepez T, Vandewoestyne M, Douterloigne K, Deforce D, et al. A novel dictionary based computer vision method for the detection of cell nuclei. PLOS ONE. 2013;8(1).

13. Janssens O, De Vylder J, Aelterman J, Verstockt S, Philips W, Van Der Straeten D, et al. Leaf segmentation and parallel phenotyping for the analysis of gene networks in plants. Signal Processing Conference, Proceedings. 2013.

14. De Vylder J, Douterloigne K, Vandenbussche F, Van Der Straeten D, Philips W. A non-rigid registration method for multispectral imaging of plants. Proceedings of SPIE, the International Society for Optical Engineering. Bellingham, WA, USA: SPIE, the International Society for Optical Engineering; 2012.

15. Douterloigne K, De Vylder J, Philips W. Dynamic subsampling for image registration speedup using the mutual information variance estimate. IEEE International Conference on Image Processing (ICIP). Piscataway, NJ, USA: IEEE; 2012. p. 2961-4.

16. De Vylder J, Vandenbussche F, Hu Y, Philips W, Van Der Straeten D. Rosette Tracker: an open source image analysis tool for automatic quantification of genotype effects. PLANT PHYSIOLOGY. 2012;160(3):1149-59.

17. De Vylder J, Vandenbussche F, Hu Y, Philips W, Van Der Straeten D. Rosette Tracker: an open source image analysis tool for automatic quantification of genotype effects. PhenoDays, Abstracts. 2012. p. 31-31.

18. Aelterman J, Goossens B, Luong Q, De Vylder J, Pizurica A, Philips W. Combined non-local and multi-resolution sparsity prior in image restoration. IEEE International Conference on Image Processing, Proceedings. IEEE; 2012. p. 3049-52.

19. De Vylder J, Aelterman J, Babin D, Vandewoestyne M, Lepez T, Deforce D, et al. A novel dictionary based detection method for cell nuclei. 34th Annual International Conference of the Engineering in Medicine and Biology Society, Abstracts. 2012.

20. De Vylder J, Aelterman J, Vandewoestyne M, Lepez T, Deforce D, Philips W. Cell nuclei detection using globally optimal active contours with shape prior. Lecture notes in computer science. 2012.

21. De Vylder J, Philips W. Computational efficient segmentation of cell nuclei in 2D and 3D fluorescent micrographs. Proceedings of SPIE, the International Society for Optical Engineering. Bellingham, WA, USA: SPIE, the International Society for Optical Engineering; 2011.

22. De Vylder J, Philips W. Non-rigid registration using mutual information minimizing splines. IEEE EMBS Benelux Chapter, Annual symposium, Abstracts. IEEE Engineering in Medicine and Biology Society (EMBS). Benelux Chapter; 2011.

23. De Vylder J, De Vos W, Philips W, Manders E. 2D mapping of strongly deformable cell nuclei, based on contour matching. CYTOMETRY PART A. 2011;79A(7):580-8.

24. De Vylder J, Rooms F, Philips W. Convex formulation and global optimization for multimodal active contour segmentation. Image and Signal Processing and Analysis, 7th International symposium, Proceedings. IEEE Signal PRocessing Society; 2011.

25. De Vylder J, Aelterman J, Philips W. Robust active contour segmentation with an efficient global optimizer. Lecture notes in computer science. Berlin, Germany: Springer; 2011. p. 195-206.

26. De Vylder J, Aelterman J, Philips W. A comparative study on the optimization of global optimum active contours. Signal Processing and Communication Systems, 5th IEEE International conference. New York, NY, USA: IEEE; 2011. p. 1-10.

27. De Vylder J, Ochoa Donoso D, Philips W, Chaerle L, Van Der Straeten D. Leaf segmentation and tracking using probabilistic parametric active contours. Lecture Notes in Computer Science. Berlin, Germany: Springer; 2011. p. 75-85.

28. De Vylder J. 2D Mapping of pathological nuclei. UGent-FirW Doctoraatssymposium, 11e. Gent: Universiteit Gent. Faculteit Ingenieurswetenschappen; 2010.

29. De Vylder J, Philips W. A computational efficient external energy for active contour segmentation using edge propagation. 2010 IEEE international conference on image processing. New York, NY, USA: IEEE; 2010. p. 661-4.

30. De Vylder J, Douterloigne K, Philips W. A fast external force field for parametric active contour segmentation. Lecture notes in computer science. Berlin, Germany: Springer; 2010. p. 125-34.

31. De Vylder J, De Vos W, Philips W. 2D mapping of cell nuclei based on contour warping. European Signal Processing Conference, 18th, Proceedings. European Association for Signal Processing (EURASIP); 2010.

32. De Vylder J, Rooms F, Philips W. Automatic detection of cell nuclei in fluorescence micrographs. Focus on Microscopy, Abstracts. 2010.

33. De Vylder J, Douterloigne K, De Vos W, Philips W. 2D Mapping of strongly deformed cell nuclei based on contour warping. IEEE Engineering in Medicine and Biology Society Conference Proceedings. New York, NY, USA: IEEE; 2010. p. 4379-82.

34. De Vylder J, Rooms F, Philips W. Segmentation of cell nuclei in Arabidopsis thaliana roots. Image analysis and recognition, 2010, PT II, proceedings. Berlin, Germany: Springer; 2010. p. 207-16.

35. De Vylder J, Ochoa Donoso D, Philips W, Chaerle L, Van Der Straeten D. Automatic analysis of biological images. Knowledge for Growth, Abstracts. 2010.

36. De Vylder J, Philips W. 2-D shape representation using improved Fourier descriptors. 2009 16TH IEEE international conference on image processing, vols 1-6. New York, NY, USA: IEEE; 2009. p. 397-400.

37. De Vylder J, Ochoa Donoso D, Philips W, Chaerle L, Van Der Straeten D. Tracking multiple objects using moving snakes. 2009 16th international conference on digital signal processing , vols 1 and 2. New York, NY, USA: IEEE; 2009. p. 966-71.

38. De Vylder J, Ochoa Donoso D, Philips W, Chaerle L, Van Der Straeten D. Tracking multiple objects using intensity-GVF snakes. European Signal Processing Conference, 17th, Proceedings. 2009.

39. De Vylder J. Improved Fourier descriptors for 2-D shape representation. UGent-FirW Doctoraatssymposium, 9e. Gent: Universiteit Gent. Faculteit Ingenieurswetenschappen; 2008.

40. De Vylder J, Philips W. Improved Fourier descriptors for 2-D shape representation. ProRISC 2008: 19th Annual workshop on circuits, systems and signal processing. 2008.

41. De Vylder J, Philips W, Gautama S, Bellens R. Characterization and recognition of cell structures in biomedical images. ProRISC 2007 18th Annual Workshop on Circuits, Systems and Signal Processing. 2007.

Bibliography

[Abbasi 97] Sadegh Abbasi, Farzin Mokhtarian & Josef Kittler. *Reliable classification of chrysanthemum leaves through curvature scale space*. In Scale-Space Theory in Computer Vision, pages 284–295. Springer, 1997.

[Abbasi 99] Sadegh Abbasi, Farzin Mokhtarian & Josef Kittler. *Curvature scale space image in shape similarity retrieval*. Multimedia systems, vol. 7, no. 6, pages 467–476, 1999.

[Aboelela 05] A. Aboelela, A. Liptay & J. L. Barron. *Plant Growth Measurement Techniques using Near-Infrared Imagery*. International Journal of Robotics and Automation, vol. 20, 2005.

[Abramoff 04] Paulo J.; Ram Sunanda J. Abramoff M.D.; Magalhaes. *Image processing with ImageJ*. Biophotonics international, vol. 11, no. 7, pages 36–42, 2004.

[Aelterman 12] J. Aelterman, B. Goossens, H. Luong, J. De Vylder, A. Pižurica & W. Philips. *Combined Non-Local and Multi-Resolution Sparsity Prior in Image Restoration*. In Proc. of the IEEE International Conference on Image Processing (ICIP2012), Orlando, Florida, USA, September 2012.

[Aelterman 13] J. Aelterman, B. Goossens, J. De Vylder, A. Pižurica & W. Philips. *Computationally efficient Locally Adaptive Demosaicing of Color Filter Array Images using the Dual-Tree Complex Wavelet Packet Transform*. PLOS ONE, 2013.

[Aelterman 14] J. Aelterman. *Multiresolution models in image restoration and reconstruction techniques with medical and other applications*. PhD thesis, Ghent university, 2014.

[Aguado 98] A. S. Aguado, M. S. Nixon & M. E. Montiel. *Parameterizing arbitrary shapes via Fourier descriptors for evidence-gathering extraction*. Computer Vision and Image Understanding, vol. 69, no. 2, pages 202–221, Feb 1998.

[Al-Kofahi 08] Yousef Al-Kofahi, Natalie Dowell-Mesfin, Christopher Pace, William Shain, James N Turner & Badrinath Roysam. *Improved detection of branching points in algorithms for automated neuron tracing from 3D confocal images*. Cytometry Part A, vol. 73, no. 1, pages 36–43, 2008.

[Al-Kofahi 10] Y. Al-Kofahi, W. Lassoued, W. Lee & B. Roysam. *Improved Automatic Detection and Segmentation of Cell Nuclei in Histopathology Images*. Biomedical Engineering, IEEE Transactions on, vol. 57, no. 4, pages 841–852, 2010.

[Aljabar 09] P. Aljabar, R.A. Heckemann, A. Hammers, J.V. Hajnal & D. Rueckert. *Multi-atlas based segmentation of brain images: Atlas selection and its effect on accuracy*. NeuroImage, vol. 46, no. 3, pages 726 – 738, 2009.

[Alsberg 11] B.K. Alsberg, T. Loke & I. Baarstad. *PryJector: A Device for In Situ Visualization of Chemical and Physical Property Distributions on Surfaces Using Projection and Hyperspectral Imaging*. Journal of Forensic Sciences, vol. 54, pages 976–983, 2011.

[Amanatiadis 11] A Amanatiadis, VG Kaburlasos, A Gasteratos & SE Papadakis. *Evaluation of shape descriptors for shape-based image retrieval*. Image Processing, IET, vol. 5, no. 5, pages 493–499, 2011.

[Amayeh 07] Gholamreza Amayeh, Shohreh Kasaei, George Bebis, Alireza Tavakkoli & Konstantinos Veropoulos. *Improvement of Zernike moment descriptors on affine transformed shapes*. In Signal Processing and Its Applications, 2007. ISSPA 2007. 9th International Symposium on, pages 1–4. IEEE, 2007.

[Arad 94] Dyn Nira Reisfeld Daniel Yeshurun Yehezkel Arad Nur. *Image warping by radial basis functions: applications to facial expressions*. 1994.

[Araus 14] JosÃ© Luis Araus & Jill E. Cairns. *Field high-throughput phenotyping: the new crop breeding frontier*. Trends in Plant Science, vol. 19, no. 1, pages 52 – 61, 2014.

[Arvidsson 11] S. Arvidsson, P. Perez-Rodriguez & B. Mueller-Roeber. *A growth phenotyping pipeline for Arabidopsis thaliana integrating image analysis and rosette area modeling for robust quantification of genotype effects*. New Phytol, vol. 191, no. 3, pages 895–907, Aug 2011.

[Aujol 05] Jean-Francois Aujol & Antonin Chambolle. *Dual Norms and Image Decomposition Models*. International Journal of Computer Vision, vol. 63, no. 1, pages 85–104, 2005.

[Aymen El Ghoul 09] Josiane Zerbia Aymen El Ghoul Ina Jermyn. *Inflection point phase field higher-order active contour model for network extraction from VHR satterlite images*. Eusipco 2009, 2009.

[Babaii Rizvandi 08] N. Babaii Rizvandi, A. Pizurica & W. Philips. *Active Appearance Model (AAM) - From theory to implementation*. In Proc. 3rd Int. Conference on Computer Vision Theory and Applications (VISAPP 2008), volume 1, pages 539–542, Funchal, Portugal, January 22-25 2008.

[Babin 12] Danilo Babin, Ewout Vansteenkiste, Aleksandra Pizurica & Wilfried Philips. *Centerline calculation for extracting abdominal aorta in 3-D MRI images*. In IEEE Engineering in Medicine and Biology Society Conference Proceedings, pages 3982–3985. IEEE, 2012.

[Babin 13a] Danilo Babin. *Segmentation and skeletonization techniques for cardiovascular image analysis*. PhD thesis, Ghent university, 2013.

[Babin 13b] Danilo Babin, Aleksandra Pizurica, Jonas De Vylder, Ewout Vansteenkiste & Wilfried Philips. *Brain blood vessel segmentation using line-shaped profiles*. PHYSICS IN MEDICINE AND BIOLOGY, vol. 58, no. 22, pages 8041–8061, 2013.

[Babin 14] Danilo Babin, M. Spyrantis, Alexandra Pižurica & Wilfried Philips. *skeleton calculation for automatic extraction of arteriovenous malformation in 3-D CTA images*. In Submitted to: 2014 IEEE International Symposium on Biomedical Imaging, 2014.

[Backes 12] André R Backes, João B Florindo & Odemir M Bruno. *Shape analysis using fractal dimension: A curvature based approach*. arXiv preprint arXiv:1201.3118, 2012.

[Baecker 07] Volker Baecker. *Automatic measurement of plant features using ImageJ and MRI Cell Image Analyzer*. Rapport technique, 2007.

[Baraniuk 07] R. G. Baraniuk. *Compressive sensing*. IEEE Signal Processing Magazine, vol. 24, no. 4, pages 118–+, Jul 2007.

[Baraniuk 10] R. G. Baraniuk, E. Candes, M. Elad & Y. Ma. *Applications of Sparse Representation and Compressive Sensing*. Proceedings of the IEEE, vol. 98, no. 6, pages 906–909, Jun 2010.

[Bardell 04] David Bardell. *The Invention of the Microscope*. BIOS, vol. 75, pages 78–84, 2004.

[Barequet 96] G. Barequet & M. Sharir. *Piecewise-linear interpolation between polygonal slices*. Computer Vision and Image Understanding, vol. 63, no. 2, pages 251–272, Mar 1996.

[Barron 94] J. L. Barron & A. Liptay. *Optical flow to measure minute increments in plant growth*. Bioimaging, vol. 2, no. 1, pages 57–61, 1994.

[Barron 97] J. L. Barron & A. Liptay. *Measuring 3-D plant growth using optical flow*. Bioimaging, vol. 5, no. 2, pages 82–86, 1997.

[Bartolini 05] I. Bartolini, P. Ciaccia & M. Patella. *WARP: Accurate retrieval of shapes using phase of Fourier descriptors and time warping distance*. IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 27, no. 1, pages 142–147, Jan 2005.

[Bayraktar 06] Bulent Bayraktar, Tytus Bernas, J. Paul Robinson & Bartek Rajwa. *Image reconstruction from discrete Chebyshev moments via formation of lookup tables*. In Proc. SPIE: Medical Imaging 2006: Physics of Medical Imaging, volume 6142, pages 61424T–61424T–7, 2006.

[Beatson 07] R. K. Beatson, M. J. D. Powell & A. M. Tan. *Fast evaluation of polyharmonic splines in three dimensions*. IMA Journal of Numerical Analysis, vol. 27, no. 3, pages 427–450, Jul 2007.

[Bergbauer 13] J. Bergbauer, C. Nieuwenhuis, M. Souiai & D. Cremers. *Proximity Priors for Variational Semantic Segmentation and Recognition*. In ICCV Workshop on Graphical Models for Scene Understanding, 2013.

[Berger 12] Bettina Berger. *High-throughput phenotyping at The Plant Accelerator*. In Phenodays, 2012.

[Bertrand 82] O. Bertrand, R. Queval & H. Maitre. *Shape Interpolation Using Fourier Descriptors with Application to Animation Graphics*. Signal Processing, vol. 4, no. 1, pages 53–58, 1982.

[Bijar 11] A. Bijar & M. Mikaeili. *Sperm's tail identification and discrimination in microscopic images of stained human semen smear*. In Image and Signal Processing and Analysis (ISPA), 2011 7th International Symposium on, pages 709–714, 2011.

[Bilmes 97] Jeff A. Bilmes. *A gentle tutorial on the EM algorithm and its application to parameter estimation for gaussian mixture and hidden markov models*. Rapport technique, International Computer Science Institute, 1997.

[Bishop 06] Christopher M Bishop*et al.* Pattern recognition and machine learning, volume 1. springer New York, 2006.

[Blake 97] A. Blake & M. Isard. *The CONDENSATION algorithm - Conditional density propagation and applications to visual tracking*. Advances in

Neural Information Processing Systems 9, vol. 9, pages 361–367 1098, 1997.

[Blake 11] Andrew Blake, Pushmeet Kohli & Carsten Rother. Markov random fields for vision and image processing. MIT Press, 2011.

[Bookstein 89] F. L. Bookstein. *Principal Warps - Thin-Plate Splines and the Decomposition of Deformations*. IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 11, no. 6, pages 567–585, Jun 1989.

[Borghgraef 10] Alexander Borghgraef, Olivier Barnich, Fabian Lapierre, Marc Van Droogenbroeck, Wilfried Philips & Marc Acheroy. *An evaluation of pixel-based methods for the detection of floating objects on the sea surface*. EURASIP journal on advances in signal processing, page 11, 2010.

[Boveiri 10] H Boveiri. *On Pattern Classification Using Statistical Moments*. International Journal of Signal Processing, Image Processing and Pattern Recognition rocessing and Pattern Recognition, vol. 3, pages 15–24, 2010.

[Boyd 04] Stephen P. Boyd & Lieven Vandenberghe. Convex optimization. Cambridge University Press, Cambridge, UK ; New York, 2004.

[Bradhurst 08] C.J. Bradhurst, W. Boles & Yin Xiao. *Segmentation of bone marrow stromal cells in phase contrast microscopy images*. In Image and Vision Computing New Zealand, 2008. IVCNZ 2008. 23rd International Conference, pages 1–6, 2008.

[Breiman 01] Leo Breiman. *Random Forests*. Machine Learning, vol. 45, no. 1, pages 5–32, 2001.

[Bresson 07a] X. Bresson & T. F. Chan. *Active contours based on Chambolle's mean curvature motion*. 2007 IEEE International Conference on Image Processing, Vols 1-7, pages 33–36 3371, 2007.

[Bresson 07b] X. Bresson, S. Esedoglu, P. Vandergheynst, J. P. Thiran & S. Osher. *Fast global minimization of the active Contour/Snake model*. Journal of Mathematical Imaging and Vision, vol. 28, no. 2, pages 151–167, Jun 2007.

[Bresson 09] Xavier Bresson. *A Short Guide on a Fast Global Minimization Algorithm for Active Contour Models*. Energy, pages 1–19, 2009.

[Brullmann 12] Dan Dominik Brullmann, Andreas Pabst, KarlM. Lehmann, Thomas Ziebart, MarcO. Klein & Bernd d'Hoedt. *Counting touching*

cell nuclei using fast ellipse detection to assess in vitro cell characteristics: a feasibility study. Clinical Oral Investigations, vol. 16, no. 1, pages 33–38, 2012.

[Campbell 11] Jamie Campbell, David Lowe & Matthew A Sleeman. *Developing the next generation of monoclonal antibodies for the treatment of rheumatoid arthritis*. British Journal of Pharmacology, vol. 162, no. 7, pages 1470–1484, 2011.

[Canny 86] John Canny. *A Computational Approach to Edge Detection*. Pattern Analysis and Machine Intelligence, IEEE Transactions on, vol. 8, pages 679–698, 1986.

[Capo-Chichi 09] C. D. Capo-Chichi, K. Q. Cai, J. R. Testa, A. K. Godwin & X. X. Xu. *Loss of GATA6 Leads to Nuclear Deformation and Aneuploidy in Ovarian Cancer*. Molecular and Cellular Biology, vol. 29, no. 17, pages 4766–4777, 2009.

[Carpenter 06] A. E. Carpenter, T. R. Jones, M. R. Lamprecht, C. Clarke, I. H. Kang, O. Friman, D. A. Guertin, J. H. Chang, R. A. Lindquist, J. Moffat, P. Golland & D. M. Sabatini. *CellProfiler: image analysis software for identifying and quantifying cell phenotypes*. Genome Biology, vol. 7, no. 10, pages –, 2006.

[Cattrysse 13] Hannes Cattrysse. Optimization of segmentation and classification of cell nuclei in the context of nuclear pathologies. Master's thesis, Ghent university, Belgium, 2013.

[Chaerle 00] L. Chaerle & D. Van Der Straeten. *Imaging techniques and the early detection of plant stress*. Trends Plant Sci, vol. 5, no. 11, pages 495–501, Nov 2000.

[Chaerle 01] Laury Chaerle & Dominique Van Der Straeten. *Seeing is believing: imaging techniques to monitor plant health*. Biochimica et Biophysica Acta (BBA)-Gene Structure and Expression, vol. 1519, no. 3, pages 153–166, 2001.

[Chaerle 04] L. Chaerle, D. Hagenbeek, E. De Bruyne, R. Valcke & D. Van der Straeten. *Thermal and chlorophyll-fluorescence imaging distinguish plant-pathogen interactions at an early stage*. Plant and Cell Physiology, vol. 45, no. 7, pages 887–896, 2004.

[Chaerle 06] L. Chaerle, D. Hagenbeek, D. Van der Straeten, I. Leinonen & H. Jones. *Monitoring and screening plant populations with thermal and*

chlorophyll fluorescence imaging. Comparative Biochemistry and Physiology a-Molecular & Integrative Physiology, vol. 143, no. 4, pages S143–S144, 2006.

[Chaerle 07] L. Chaerle, I. Leinonen, H. G. Jones & D. Van der Straeten. *Monitoring and screening plant populations with combined thermal and chlorophyll fluorescence imaging*. Journal of Experimental Botany, vol. 58, no. 4, pages 773–784, 2007.

[Chaker 07] Fatma Chaker, M Tarak Bannour & Faouzi Ghorbel. *Contour Retrieval and Matching by Affine Invariant Fourier Descriptors.* MVA, vol. 7, pages 291–294, 2007.

[Chaker 08] F. Chaker, F. Ghorbel & M. T. Bannour. *Content-based shape retrieval using different affine shape descriptors*. Visapp 2008: Proceedings of the Third International Conference on Computer Vision Theory and Applications, Vol 1, pages 497–500 230, 2008.

[Chambolle 11] A. Chambolle & T. Pock. *A First-Order Primal-Dual Algorithm for Convex Problems with Applications to Imaging*. Journal of Mathematical Imaging and Vision, vol. 40, no. 1, pages 120–145, 2011.

[Chan 99] T. Chan & L. Vese. *An active contour model without edges*. Scale-Space Theories in Computer Vision, vol. 1682, pages 141–151, 1999.

[Chan 06] T. F. Chan, S. Esedoglu & M. Nikolova. *Algorithms for finding global minimizers of image segmentation and denoising models*. Siam Journal on Applied Mathematics, vol. 66, no. 5, pages 1632–1648, 2006.

[Chard 92] T. Chard. *Pregnancy tests: a review*. Human Reproduction, vol. 7, no. 5, pages 701–710, 1992.

[Charmi 08] M. A. Charmi, S. Derrode & F. Ghorbel. *Fourier-based geometric shape prior for snakes*. Pattern Recognition Letters, vol. 29, no. 7, pages 897–904, May 1 2008.

[Chen 06] X. W. Chen, X. B. Zhou & S. T. C. Wong. *Automated segmentation, classification, and tracking of cancer cell nuclei in time-lapse microscopy*. IEEE Transactions on Biomedical Engineering, vol. 53, no. 4, pages 762–766, Apr 2006.

[Chen 10] Fuhua Chen & Yunmei Chen. *A stochastic variational model for multiphase soft segmentation with bias correction*. Advanced Modeling and Optimization. v12 i3, pages 339–345, 2010.

[Chen 12] Dai-Qiang Chen, Hui Zhang & Li-Zhi Cheng. *A Fast Fixed Point Algorithm for Total Variation Deblurring and Segmentation*. Journal of Mathematical Imaging and Vision, pages 1–13, 2012.

[Chinta 12] Rambabu Chinta & Martin Wasser. *Three-dimensional segmentation of nuclei and mitotic chromosomes for the study of cell divisions in live Drosophila embryos*. Cytometry Part A, vol. 81A, no. 1, pages 52–64, 2012.

[Chitwood 14] Daniel H. Chitwood, Aashish Ranjan, Ciera C. Martinez, Lauren R. Headland, Thinh Thiem, Ravi Kumar, Michael F. Covington, Tommy Hatcher, Daniel T. Naylor, Sharon Zimmerman, Nora Downs, Nataly Raymundo, Edward S. Buckler, Julin N. Maloof, Mallikarjuna Aradhya, Bernard Prins, Lin Li, Sean Myles & Neelima R. Sinha. *A Modern Ampelography: A Genetic Basis for Leaf Shape and Venation Patterning in Grape*. Plant Physiology, vol. 164, no. 1, pages 259–272, 2014.

[Chong 03a] C-W Chong, P Raveendran & R Mukundan. *The scale invariants of pseudo-Zernike moments*. Pattern Analysis & Applications, vol. 6, no. 3, pages 176–184, 2003.

[Chong 03b] Chee-Way Chong, P Raveendran & R Mukundan. *Translation invariants of Zernike moments*. Pattern Recognition, vol. 36, no. 8, pages 1765–1773, 2003.

[Clifford 90] Peter Clifford. *Markov random fields in statistics*. Disorder in physical systems, pages 19–32, 1990.

[Cloppet 08a] F. Cloppet & A. Boucher. *Segmentation of overlapping/aggregating nuclei cells in biological images*. In 19th International Conference on Pattern Recognition, Vols 1-6, pages 789–792, 2008.

[Cloppet 08b] F. Cloppet & A. Boucher. *Segmentation of overlapping/aggregating nuclei cells in biological images*. 19th International Conference on Pattern Recognition, Vols 1-6, pages 789–792, 2008.

[Cloppet 10] F. Cloppet & A. Boucher. *Segmentation of complex nucleus configurations in biological images*. Pattern Recognition Letters, vol. In Press, Corrected Proof, 2010.

[Cofaru 12] Corneliu Cofaru. *Novel Digital Image Correlation Approaches for the Measurement of Heterogeneous Displacements and Strains*. PhD thesis, Ghent University, 2012.

[Cohen 93] L. D. Cohen & I. Cohen. *Finite-Element Methods for Active Contour Models and Balloons for 2-D and 3-D Images*. IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 15, no. 11, pages 1131–1147, Nov 1993.

[Cootes 95] T. F. Cootes, C. J. Taylor, D. H. Cooper & J. Graham. *Active Shape Models - Their Training and Application*. Computer Vision and Image Understanding, vol. 61, no. 1, pages 38–59, Jan 1995.

[Cootes 01] T. F. Cootes & C. J. Taylor. *Statistical models of appearance for medical image analysis and computer vision*. Medical Imaging: 2001: Image Processing, Pts 1-3, vol. 2, no. 27, pages 236–248 1968, 2001.

[Cote 07] D. Cote, J.A. Spencer, D.P. Biss, C.P. Lin & I. Veilleux. *Multimodal nonlinear and linear video rate microscopy in neuroscience*. In 11th World MultiConference on Systemics Cybernetics and Informatics Vol V Post Conference Issue Proceedings, 2007.

[Cremers 08] D. Cremers, F. R. Schmidt & F. Barthel. *Shape priors in variational image segmentation: Convexity, Lipschitz continuity and globally optimal solutions*. 2008 IEEE Conference on Computer Vision and Pattern Recognition, Vols 1-12, pages 810–815, 2008.

[Crimmins 82] T. R. Crimmins. *A Complete Set of Fourier Descriptors for Two-Dimensional Shapes*. IEEE Transactions on Systems Man and Cybernetics, vol. 12, no. 6, pages 848–855, 1982.

[Dambreville 06] Samuel Dambreville, Yogesh Rathi & Allen Tannenbaum. *Tracking deformable objects with unscented Kalman filtering and geometric active contours*. In American Control Conference, 2006, pages 6–pp. IEEE, 2006.

[Davidson 02] Michael W Davidson & Mortimer Abramowitz. *Optical microscopy*. Encyclopedia of imaging science and technology, 2002.

[De Bock 05] J. De Bock, P. De Smet & W. Philips. *Image segmentation using watersheds and normalized cuts*. Proceedings of SPIE, Vision Geometry XIII, vol. 5675, pages 164–173, 2005.

[De Bock 06a] J. De Bock, R. Pires, P. De Smet & W. Philips. *A fast dynamic border linking algorithm for region merging*. In Advanced Concepts for Intelligent Vision Systems, ACIVS, 8th international conference, Lecture Notes in Computer Science, pages 232–241, 2006.

[De Bock 06b] J. De Bock, R. Pires, P. De Smet & W. Philips. *A fast dynamic border linking algorithm for region merging*. Advanced Concepts for Intelligent Vision Systems, ACIVS, 8th international conference, pages 232–241, 2006.

[De Bock 07] J. De Bock & W. Philips. *Line segment based watershed segmentation*. In Advances in Computer Vision and Computer Graphics, Third International Conference on Computer Vision and Computer Graphics, Lecture Notes in Computer Science, pages 579–586, 2007.

[De Bock 10] Johan De Bock & Wilfried Philips. *Fast and memory efficient 2-D connected components using linked lists of line segments*. IEEE TRANSACTIONS ON IMAGE PROCESSING, vol. 19, no. 12, pages 3222–3231, 2010.

[De Smet 00] P. De Smet & R. Pires. *Implementation and Analysis of an Optimized Rainfalling Watershed Algorithm*. In proc. of Image and Video Communications and Processing 2000. SPIE, 2000.

[De Vos 09a] W. H. De Vos, R. A. Hoebe, G. H. Joss, W. Haffinans, S. Baatout, P. Van Oostveldt & E. M. M. Manders. *Controlled Light Exposure Microscopy Reveals Dynamic Telomere Microterritories Throughout the Cell Cycle*. Cytometry Part A, vol. 75A, no. 5, pages 428–439, May 2009.

[De Vos 09b] W. H. De Vos, R. A. Hoebe, G. H. Joss, W. Haffmans, S. Baatout, P. Van Oostveldt & E. M. Manders. *Controlled light exposure microscopy reveals dynamic telomere microterritories throughout the cell cycle*. Cytometry A, vol. 75, no. 5, pages 428–39, May 2009.

[De Vos 10a] W. H. De Vos, F. Houben, R. A. Hoebe, R. Hennekam, B. van Engelen, E. M.M. Manders, F. C.S. Ramaekers, J. L.V. Broers & Van Oostveldt P. *Increased plasticity of the nuclear envelope and hypermobility of telomeres due to the loss of A-type lamins*. Biochimica et Biophysica Acta (BBA) - General Subjects, vol. In Press, Corrected Proof, pages 448–458, 2010.

[De Vos 10b] W. H. De Vos, L. Van Neste, B. Dieriks, G. H. Joss & P. Van Oostveldt. *High Content Image Cytometry in the Context of Subnuclear Organization*. Cytometry Part A, vol. 77A, no. 1, pages 64–75, Jan 2010.

[De Vylder 07] J. De Vylder, R. Bellens, S. Gautama & W. Philips. *Characterization and recognition of cell structures in biomedical images*. In IEEEPRORISC Workshop on Circuits, Systems and Signal Processing, November 2007.

[De Vylder 08a] J. De Vylder & W. Philips. *Improved Fourier Descriptors for 2-D Shape Representation*. In IEEE PRORISC Workshop on Circuits, Systems and Signal Processing, November 2008.

[De Vylder 08b] Jonas De Vylder. *Improved Fourier descriptors for 2-D shape representation*. In UGent-FirW Doctoraatssymposium, 9e, page 2. Universiteit Gent. Faculteit Ingenieurswetenschappen, 2008.

[De Vylder 09a] J. De Vylder, D. Ochoa, W. Philips, L. Chaerle & D. Van Der Straeten. *Tracking Multiple Objects Using Intensity-GVF Snakes*. In 17th European Signal Processing Conference 2009 (EUSIPCO 2009), Glasgow, Scotland, aug 2009.

[De Vylder 09b] J. De Vylder, D. Ochoa, W. Philips, L. Chaerle & D. Van Der Straeten. *Tracking Multiple Objects Using Moving Snakes*. In 2009 16th International Conference on Digital Signal Processing, pages 966–971, 2009.

[De Vylder 09c] J. De Vylder & W. Philips. *2-D Shape Representation Using Improved Fourier Descriptors*. In 2009 16th IEEE International Conference on Image Processing, pages 397–400, 2009.

[De Vylder 10a] J. De Vylder, W. De Vos & W. Philips. *2D mapping of cell nuclei based on contour warping*. In European Signal Processing Conference (EUSIPCO 2010), aug 2010.

[De Vylder 10b] J. De Vylder, W. De Vos & W. Philips. *2D Mapping of pathological nuclei*. In Proc. 1th FirW Ph.D. Symposium, Ghent, Belgium, Dec., 2010.

[De Vylder 10c] J. De Vylder, K. Douterloigne, W. De Vos & W. Philips. *2D Mapping of Strongly Deformed Cell Nuclei Based on Contour Warping*. In 2010 Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), pages 4379–4382, 2010.

[De Vylder 10d] J. De Vylder, D. Ochoa, W. Philips, L. Chaerle & D. Van Der Straeten. *Automatic analysis of biological images*. In Knowledge for Growth, May 2010.

[De Vylder 10e] J. De Vylder & W. Philips. *A Computational Efficient External Energy for Active Contour Segmentation Using Edge Propagation*. In 2010 IEEE International Conference on Image Processing, pages 661–664, 2010.

[De Vylder 10f] J. De Vylder, F. Rooms & W. Philips. *Automatic Detection of Cell Nuclei in Fluorescence Micrographs*. In Focus on Microscopy 2010, Shanghai, China, March 2010.

[De Vylder 10g] J. De Vylder, F. Rooms & W. Philips. *Segmentation of cell nuclei in Arabidopsis Thaliana Roots*. In Proc. of the 7th Int. Conf. on Image Analysis and Recognition (ICIAR2010), Lecture Notes in Computer Science, Póvoa de Varzim, Portugal, June 2010. Springer Verlag.

[De Vylder 10h] Jonas De Vylder, Koen Douterloigne & Wilfried Philips. *A Fast External Force Field for Parametric Active Contour Segmentation*. In Advanced Concepts for Intelligent Vision Systems, volume 6474, pages 125–134. Springer Berlin / Heidelberg, 2010.

[De Vylder 11a] J. De Vylder, J. Aelterman & W. Philips. *A comparative study on the optimization of global optimum active contours*. In The 5th IEEE International Conference on Signal Processing and Communication Systems, Honolulu, Hawaii, USA, December 2011. IEEE Communications Society, IEEE Communications Society.

[De Vylder 11b] J. De Vylder, J. Aelterman & W. Philips. *Robust Active Contour Segmentation with an Efficient Global Optimizer*. In Proc. Advanced Concepts for Intelligent Vision Systems (ACIVS) 2011, pages 195–206, Ghent, Belgium, August 2011.

[De Vylder 11c] J. De Vylder, W. H. De Vos, E. M. Manders & W. Philips. *2D Mapping of Strongly Deformable Cell Nuclei, Based on Contour Matching*. Cytometry Part A, vol. 79A, no. 7, pages 580–588, Jul 2011.

[De Vylder 11d] J. De Vylder, D. Ochoa, L. Chaerle, D. Van Der Straeten & W. Philips. *Leaf Segmentation and Tracking Using Probabilistic Parametric Active Contours*. In Proc. of Mirage 2011, Lecture Notes in Computer Science. Springer Verlag, October 2011.

[De Vylder 11e] J. De Vylder & W. Philips. *Computational efficient segmentation of cell nuclei in 2D and 3D fluorescent micrographs*. In Imaging, Manipulation, and Analysis of Biomolecules, Cells, and Tissues IX, proceedings of SPIE Vol. 7902, 2011.

[De Vylder 11f] J. De Vylder & W. Philips. *non-rigid registration using mutual information minimizing splines*. In the Annual IEEE EMBS Benelux Symposium, December 2011.

[De Vylder 11g] J. De Vylder, F. Rooms & W. Philips. *Convex formulation and global optimization for multimodal active contour segmentation*. In Proc.

of International Symposium on Image and Signal Processing and Analysis 2011, Dubrovnik, Croatia, September 2011. IEEE Signal Processing Society, IEEE Signal Processing Society.

[De Vylder 12a] J. De Vylder, J. Aelterman, M. Vandewoestyne, T. Lepez, D. Deforce & W. Philips. *Cell nuclei detection using globally optimal active contours with shape prior*. In George Bebis, Richard Boyle, Bahram Parvin, Darko Koracin, Charless Fowlkes, Sen Wang, Min-Hyung Choi, Stephan Mantler, Jürgen Schulze, Daniel Acevedo, Klaus Mueller & Michael Papka, editeurs, Advances in Visual Computing, volume 7432 of *Lecture Notes in Computer Science*, pages 115–124, Crete, Greece, July 2012. Springer.

[De Vylder 12b] J. De Vylder, K. Douterloigne, F. Vandenbussche, D. Van Der Straeten & W. Philips. *A non-rigid registration method for multispectral imaging of plants*. In SPIE Proc. of Sensing for Agriculture and Food Quality and Safety, Baltimore, Maryland, USA, April 2012. SPIE.

[De Vylder 12c] J. De Vylder, F. Vandenbussche, Y. Hu, W. Philips & D. Van Der Straeten. *Rosette Tracker: an open source image analysis tool for automatic quantification of genotype effects*. Plant Physiology, 2012.

[De Vylder 12d] J. De Vylder, F. Vandenbussche, Y. Hu, W. Philips & D. Van Der Straeten. *Rosette Tracker: an open source image analysis tool for automatic quantification of genotype effects*. In Phenodays 2012, October 2012.

[De Vylder 12e] Jonas De Vylder, Jan Aelterman, Danilo Babin, Mado Vandewoestyne, Trees Lepez, Dieter Deforce & Wilfried Philips. *A novel dictionary based detection method for cell nuclei*. In 34th Annual International Conference of the Engineering in Medicine and Biology Society, Abstracts, 2012.

[De Vylder 13a] J. De Vylder, J. Aelterman, M. Vandewoestyne, T. Lepez, K. Douterloigne, D. Deforce & W. Philips. *A Novel Dictionary Based Computer Vision Method for the Detection of Cell Nuclei*. PLOS ONE, vol. 8, no. e54068, 2013.

[De Vylder 13b] J. De Vylder, W. Philips & D. Van Der Straeten. *Multiple Leaf Tracking Using Computer Vision Methods with Shape Constraints*. In Proc. of the International Conference on Sensing Technologies for Biomaterial, Food, and Agriculture (SeTBio2013), Yokohoma, Japan, April 2013.

[De Vylder 13c] J. De Vylder, F. Rooms, S. Dhondt, D. Inze & W. Philips. *A novel tracing method for the segmentation of cell wall networks*. In Proc. of Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC13), Osaka, Japan, July 2013.

[De Vylder 14a] J. De Vylder & W. Philips. *Image segmentation, yet another inverse problem*. IEEE Signal Processing Magazine, 2014.

[De Vylder 14b] J. De Vylder & W. Philips. *Improved Fourier transform: shape description using the non-uniform Fourier transform*. Plos One, page in preparation, 2014.

[De Vylder 14c] J. De Vylder, D. Van Haerenborgh, J. Aelterman & W. Philips. *A Probabilistic Interpretation of Geometric Active Contour Segmentation*. In 2014 European Signal Processing Conference (EUSIPCO 2014), 2014.

[Deboeverie 10] F. Deboeverie, K. Teelen, P. Veelaert & W. Philips. *Adaptive Constructive Polynomial Fitting*. In Proc. of Advanced Concepts for Intelligent Vision Systems 2010(ACIVS2010), Lecture Notes in Computer Science, pages 173–184, Sydney, Australia, December 2010. Springer Verlag.

[Deboeverie 11] F. Deboeverie, P. Veelaert & W. Philips. *Face analysis using curve edge maps*. In Guiseppe Maino & Gian Luca Foresti, editeurs, Lecture Notes in Computer Science, volume 6979, pages 109–118. Springer, 2011.

[Deboeverie 13a] F. Deboeverie, P. Veelaert & W. Philips. *Image segmentation with adaptive region growing based on a polynomial surface model*. Journal of Electronic Imaging, vol. 22, no. 4, pages 043004–043004, 2013.

[Deboeverie 13b] F. Deboeverie, P. Veelaert & W. Philips. *Segmented face approximation with adaptive region growing based on low-degree polynomial fitting*. Springer Journal on Signal, Image and Video Processing (SIViP), Mar 2013.

[Demanet 07] Laurent Demanet & Lexing Ying. *Wave atoms and sparsity of oscillatory patterns*. Applied and Computational Harmonic Analysis, vol. 23, no. 3, pages 368–387, 2007.

[Denk 04] Winfried Denk & Heinz Horstmann. *Serial block-face scanning electron microscopy to reconstruct three-dimensional tissue nanostructure*. PLoS biology, vol. 2, no. 11, page e329, 2004.

[Derrode 06] S. Derrode, M. A. Charmi & F. Ghorbel. *Fourier-based invariant shape prior for snakes*. 2006 IEEE International Conference on Acoustics, Speech and Signal Processing, Vols 1-13, pages 1349–1352 6086, 2006.

[Despotovic 14] Ivana Despotovic. *Multimodal Image Analysis of the Human Brain*. PhD thesis, Ghent university, 2014.

[Dhondt 12a] Stijn Dhondt. *Imaging as a tool to study leaf development in Arabidopsis thaliana*. PhD thesis, Ghent university, 2012.

[Dhondt 12b] Stijn Dhondt, Dirk Van Haerenborgh, Caroline Van Cauwenbergh, Roeland M. H. Merks, Wilfried Philips, Gerrit T. S. Beemster & Dirk Inze. *Quantitative analysis of venation patterns of Arabidopsis leaves by supervised image analysis*. The Plant Journal, vol. 69, no. 3, pages 553–563, 2012.

[Di Ruberto 02] C. Di Ruberto, A. Dempster, S. Khan & B. Jarra. *Analysis of infected blood cell images using morphological operators*. Image and Vision Computing, vol. 20, no. 2, pages 133–146, 2002.

[Dice 45] Lee R. Dice. *Measures of the Amount of Ecologic Association Between Species*. Ecology, vol. 26, no. 3, pages pp. 297–302, 1945.

[Dieriks 09] Birger Dieriks, Winnok De Vos, Geert Meesen, Kaat Van Oostveldt, Tim De Meyer, Myriam Ghardi, Sarah Baatout & Patrick Van Oostveldt. *High Content Analysis of Human Fibroblast Cell Cultures after Exposure to Space Radiation*. Radiation Research, vol. 172, no. 4, pages 423–436, September 2009.

[Dieriks 10] B. Dieriks, W. H. De Vos, H. Derradji, S. Baatout & P. Van Oostveldt. *Medium-mediated DNA repair response after ionizing radiation is correlated with the increase of specific cytokines in human fibroblasts*. Mutat Res, Jan 18 2010.

[Ding 10] Jian-Jiun Ding, Wei-Lun Chao, Jiun-De Huang & Cheng-Jin Kuo. *Asymmetric fourier descriptor of non-closed segments*. In Image Processing (ICIP), 2010 17th IEEE International Conference on, pages 1613–1616. IEEE, 2010.

[Dornbusch 12] Tino Dornbusch & Olivier Michaud. *Phenotyping of individual Arabidopsis leaves using the Scanalyzer HTS*. In phenodays, 2012.

[Dowson-Day 99] M. J. Dowson-Day & A. J. Millar. *Circadian dysfunction causes aberrant hypocotyl elongation patterns in Arabidopsis*. Plant J, vol. 17, no. 1, pages 63–71, 1999.

[Du 07] Ji-Xiang Du, Xiao-Feng Wang & Guo-Jun Zhang. *Leaf shape based plant species recognition*. Applied Mathematics and Computation, vol. 185, no. 2, pages 883 – 893, 2007.

[Duan 08] Weisheng Duan, Falko Kuester, Jean-Luc Gaudiot & Omar Hammami. *Automatic object and image alignment using Fourier descriptors*. Image and Vision Computing, vol. 26, no. 9, pages 1196–1206, 2008.

[Duda 12] Richard O Duda, Peter E Hart & David G Stork. Pattern classification. John Wiley & Sons, 2012.

[Eberly 94] D. Eberly, R. Gardner, B. Morse, S. Pizer & C. Scharlach. *Ridges for image analysis*. Journal of Mathematical Imaging and Vision, vol. 4, no. 4, pages 353–373, 1994.

[El-ghazal 09] Akrem El-ghazal, Otman Basir & Saeid Belkasim. *Farthest point distance: A new shape signature for Fourier descriptors*. Signal Processing: Image Communication, vol. 24, no. 7, pages 572–586, 2009.

[Elhabian 08] Shireen Y Elhabian, Khaled M El-Sayed & Sumaya H Ahmed. *Moving object detection in spatial domain using background removal techniques-state-of-art*. Recent patents on computer science, vol. 1, no. 1, pages 32–54, 2008.

[Evans 92] L. C. Evans & R. F. Gariepy. Measure theory and fine properties of functions. Studies in advanced mathematics. CRC Press, Boca Raton, 1992.

[Felzenszwalb 04] Pedro F Felzenszwalb & Daniel P Huttenlocher. *Efficient graph-based image segmentation*. International Journal of Computer Vision, vol. 59, no. 2, pages 167–181, 2004.

[Finkelstein 94] Ruth R. Finkelstein. *Mutations at two new Arabidopsis ABA response loci are similar to the abi3 mutations*. The Plant Journal, vol. 5, no. 6, pages 765–771, 1994.

[Fischler 81] Martin A. Fischler & Robert C. Bolles. *Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography*. Commun. ACM, vol. 24, no. 6, pages 381–395, June 1981.

[Florindo 12] João Batista Florindo, André Ricardo Backes, Mário de Castro & Odemir Martinez Bruno. *A comparative study on multiscale fractal dimension descriptors*. Pattern Recognition Letters, vol. 33, no. 6, pages 798–806, 2012.

[Flusser 93a] Jan Flusser & Tomás Suk. *Character Recognition by Affine Moment Invariants*. In CAIP, pages 572–577, 1993.

[Flusser 93b] Jan Flusser & Tomás Suk. *Pattern recognition by affine moment invariants*. Pattern Recognition, vol. 26, no. 1, pages 167–174, 1993.

[Folkers 02] A. Folkers & H. Samet. *Content-based image retrieval using Fourier descriptors on a logo database*. 16th International Conference on Pattern Recognition, Vol Iii, Proceedings, pages 521–524 1068, 2002.

[Frangi 98] Ro F. Frangi, Wiro J. Niessen, Koen L. Vincken & Max A. Viergever. *Multiscale vessel enhancement filtering*. In Medical Image Computing and Computer-Assisted Interventation, lecture notes in computer science, pages 130–137. Springer-Verlag, 1998.

[Frangi 99] A. F. Frangi, W. J. Niessen, R. M. Hoogeveen, T. van Walsum & M. A. Viergever. *Model-based quantitation of 3-D magnetic resonance angiographic images*. IEEE Transactions on Medical Imaging, vol. 18, no. 10, pages 946–956, 1999.

[Freedman 05] D. Freedman, R. J. Radke, T. Zhang, Y. Jeong, D. M. Lovelock & G. T. Y. Chen. *Model-based segmentation of medical imagery by matching distributions*. IEEE Transactions on Medical Imaging, vol. 24, no. 3, pages 281–292, Mar 2005.

[French 09] A. French, S. Ubeda-Tomas, T. J. Holman, M. J. Bennett & T. Pridmore. *High-Throughput Quantification of Root Growth Using a Novel Image-Analysis Tool*. Plant Physiology, vol. 150, no. 4, pages 1784–1795, 2009.

[Gautama 03] Sidharta Gautama. *Aspecten van objectherkenning in beelden*. PhD thesis, Ghent University. Faculty of Engineering and Architecture, 2003.

[Gerlich 01] D. Gerlich, J. Beaudouin, M. Gebhard, J. Ellenberg & R. Eils. *Four-dimensional imaging and quantitative reconstruction to analyse complex spatiotemporal processes in live cells*. Nature Cell Biology, vol. 3, no. 9, pages 852–855, Sep 2001.

[Giannekou 08] Y. Giannekou, P. Tzouveli, Y. Avrithis & S. Kollias. *Affine invariant curve matching using normalization and curvature scale-space*. 2008 International Workshop on Content-Based Multimedia Indexing, pages 192–199 575, 2008.

[Giuliano 09] K. A. Giuliano, D. R. Premkumar, C. J. Strock, P. Johnston & D. L. Taylor. *Cellular Systems Biology Profiling Applied to Cellular Models of Disease*. Combinatorial Chemistry & High Throughput Screening, vol. 12, no. 9, pages 838–848, November 2009.

[Gladilin 07] E Gladilin, A Micoulet, B Hosseini, K Rohr, J Spatz & R Eils. *3D finite element analysis of uniaxial cell stretching: from image to insight*. Physical Biology, vol. 4, no. 2, page 104, 2007.

[Gladilin 08] E. Gladilin, S. Goetze, J. Mateos-Langerak, R. Van Driel, R. Eils & K. Rohr. *Shape normalization of 3D cell nuclei using elastic spherical mapping*. Journal of Microscopy-Oxford, vol. 231, no. 1, pages 105–114, Jul 2008.

[Goldenberg 01] R. Goldenberg, R. Kimmel, E. Rivlin & M. Rudzsky. *Fast geodesic active contours*. IEEE Transactions on Image Processing, vol. 10, no. 10, pages 1467–1475, Oct 2001.

[Goldstein 09] T. Goldstein & S. Osher. *The Split Bregman Method for L1-Regularized Problems*. Siam Journal on Imaging Sciences, vol. 2, no. 2, pages 323–343, 2009.

[Goldstein 10] T. Goldstein, X. Bresson & S. Osher. *Geometric Applications of the Split Bregman Method: Segmentation and Surface Reconstruction*. Journal of Scientific Computing, vol. 45, no. 1-3, pages 272–293, Oct 2010.

[Gong 04] L. X. Gong, S. D. Pathak, D. R. Haynor, P. S. Cho & Y. Kim. *Parametric shape modeling using deformable superellipses for prostate segmentation*. IEEE Transactions on Medical Imaging, vol. 23, no. 3, pages 340–349, Mar 2004.

[Gong 05] L. X. Gong, L. Ng, S. D. Pathak, I. Tutar, P. S. Cho, D. R. Haynor & Y. Kim. *Prostate ultrasound image segmentation using level set-based region flow with shape guidance*. Medical Imaging 2005: Image Processing, Pt 1-3, vol. 5747, pages 1648–1657 2200, 2005.

[Goobic 01] A. P. Goobic, M. E. Welser, S. T. Acton & K. Ley. *Biomedical application of target tracking in clutter*. Conference Record of the Thirty-Fifth Asilomar Conference on Signals, Systems and Computers, Vols 1 and 2, pages 88–92 1767, 2001.

[Goossens 10] B. Goossens. *Multresolution Image Models and Estimation Techniques*. PhD thesis, Ghent University, 2010.

[Goossens 11] B. Goossens, J. Aelterman, H. Luong, A. Pižurica & W. Philips. Discrete wavelet transforms - algorithms and applications, chapitre Wavelet-Based Analysis and Estimation of Colored Noise, pages 255–280. Intech, 2011.

[Goossens 13] Bart Goossens, Jan Aelterman, Quang Luong, Aleksandra Pizurica & Wilfried Philips. *A split-augmented Lagrangian algorithm for spectral factorization of a set of 2D directional filters and application to the design of compact shearlet frames*. In Dimitri Van De Ville, Vivek K. Goyal & Manos Papadakis, editeurs, Proc. of SPIE, volume 8858, pages 1–12. SPIE, 2013.

[Granier 06] C. Granier, L. Aguirrezabal, K. Chenu, S. J. Cookson, M. Dauzat, P. Hamard, J. J. Thioux, G. Rolland, S. Bouchier-Combaud, A. Lebaudy, B. Muller, T. Simonneau & F. Tardieu. *PHENOPSIS, an automated platform for reproducible phenotyping of plant responses to soil water deficit in Arabidopsis thaliana permitted the identification of an accession with low sensitivity to soil water deficit*. New Phytologist, vol. 169, no. 3, pages 623–635, 2006.

[Gudla 08] P. R. Gudla, K. Nandy, J. Collins, K. J. Meaburn, T. Misteli & S. J. Lockett. *A high-throughput system for segmenting nuclei using multiscale techniques*. Cytometry Part A, vol. 73A, no. 5, pages 451–466, May 2008.

[Guyer 86] DE Guyer, GE Miles, MM Schreiber, OR Mitchell & VC Vanderbilt. *Machine vision and image processing for plant identification*. Trans. ASABE, vol. 29, pages 1500–1507, 1986.

[Hall 09] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann & I. H. Witten. *The WEKA Data Mining Software: An Update*. SIGKDD Explorations, vol. 11, 2009.

[Han 07] X Han, C Xu & JL Prince. *Fast numerical scheme for gradient vector flow computation using a multigrid method*. Image Processing, IET, vol. 1, no. 1, pages 48–55, 2007.

[Harder 09] N. Harder, F. Mora-Bermudez, A. Wunsche, R. Eils, J. Ellenberg & K. Rohr. *Automatic analysis of dividing cells in live cell movies to detect mitotic delays and correlate phenotypes in time*. Genome Research, vol. 19, pages 2113–2124, 2009.

[Hodneland 13] Erlend Hodneland, Tanja Kogel, Dominik Frei, Hans-Hermann Gerdes & Arvid Lundervold. *CellSegm - a MATLAB toolbox for*

high-throughput 3D cell segmentation. Source Code for Biology and Medicine, vol. 8, no. 1, page 16, 2013.

[Hoebe 07] R. A. Hoebe, C. H. Van Oven, Jr. Gadella T. W., P. B. Dhonukshe, C. J. Van Noorden & E. M. Manders. *Controlled light-exposure microscopy reduces photobleaching and phototoxicity in fluorescence live-cell imaging*. Nat Biotechnol, vol. 25, no. 2, pages 249–53, Feb 2007.

[Horvath 06] P. Horvath, I. H. Jermyn, Z. Kato & J. Zerubia. *An improved 'gas of circles' higher-order active contour model and its application to tree crown extraction*. Computer Vision, Graphics and Image Processing, Proceedings, vol. 4338, pages 152–161, 2006.

[Horvath 09] P. Horvath, I. H. Jermyn, Z. Kato & J. Zerubia. *A higher-order active contour model of a 'gas of circles' and its application to tree crown extraction*. Pattern Recognition, vol. 42, no. 5, pages 699–709, 2009.

[Houhou 08] Nawal Houhou, J Thiran & Xavier Bresson. *Fast texture segmentation model based on the shape operator and active contour*. In Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on, pages 1–8. IEEE, 2008.

[Houhou 09] Nawal Houhou, Jean-Philippe Thiran & Xavier Bresson. *Fast Texture Segmentation Based on Semi-Local Region Descriptor and Active Contour.* Numerical Mathematics: Theory, Methods & Applications, vol. 2, no. 4, 2009.

[Hu 62] M Hu. *visual-pattern recognition by moment invariants*. IRE transactions on information theory, vol. 8, no. 2, pages 179–187, 1962.

[Hu 13] Yong Hu & Zuoyong Li. *An Improved Shape Signature for Shape Representation and Image Retrieval*. Journal of Software, vol. 8, no. 11, 2013.

[Hukkanen 10] J. Hukkanen, A. Hategan, E. Sabo & I. Tabus. *Segmentation of Cell Nuclei from Histological Images by Ellipse Fitting*. In The 2010 European Signal Processing Conference, 2010.

[Huygens 11] Huygens. *Huygens Professional: User Guide for version 4.0*. Rapport technique, Scientific Volume Imaging B.V., 2011.

[Isard 98a] M. Isard & A. Blake. Active contours. Springer, 1998.

[Isard 98b] M. Isard & A. Blake. *CONDENSATION - Conditional density propagation for visual tracking*. International Journal of Computer Vision, vol. 29, no. 1, pages 5–28, Aug 1998.

[Ishiguro 05] M. Ishiguro, I. Murase, N. Moriyama & R. Sekiguchi. *A classification method of liver tumors based on temporal change of hounsfield unit in CT images*. Medical Imaging 2005: Image Processing, Pt 1-3, vol. 5747, pages 822–830 2200, 2005.

[Iwata 02] H Iwata & Y Ukai. *SHAPE: a computer program package for quantitative evaluation of biological shapes based on elliptic Fourier descriptors*. Journal of Heredity, vol. 93, no. 5, pages 384–385, 2002.

[Jacobson 97] DL Jacobson, S J Gange, Gange SJ, N R Rose, Rose NR, N M Graham & Graham NM. *Epidemiology and estimated population burden of selected autoimmune diseases in the United States.* Clin Immunol Immunopathol., vol. 84, no. 0090-1229 (Linking), pages 223–43, 1997.

[Jaffe 85] M. J. Jaffe, A. H. Wakefield, F. Telewski, E. Gulley & R. Biro. *Computer-Assisted Image-Analysis of Plant-Growth, Thigmomorphogenesis, and Gravitropism*. Plant Physiology, vol. 77, no. 3, pages 722–730, 1985.

[Jalba 04] Andrei C Jalba, Michael HF Wilkinson & Jos BTM Roerdink. *CPM: A deformable model for shape recovery and segmentation based on charged particles*. Pattern Analysis and Machine Intelligence, IEEE Transactions on, vol. 26, no. 10, pages 1320–1335, 2004.

[Jansen 09] M. Jansen, F. Gilmer, B. Biskup, K.A. Nagel, U. Rascher, A. Fischbach, S. Briem, G. Dreissen, S. Tittmann, S. Braun, I. De Jaeger, M. Metzlaff, U. Schurr, H. Scharr & A. Walter. *Simultaneous phenotyping of leaf growth and chlorophyll fluorescence via GROWSCREEN FLUORO allows detection of stress tolerance in Arabidopsis thaliana and other rosette plants*. Functional Plant Biology, vol. 36, no. 11, pages 902–914, 2009.

[Janssens 13] O. Janssens, J. De Vylder, J. Aelterman, S. Verstockt, W. Philips, D. Van Der Straeten, S. Van Hoecke & R. Van de Walle. *Leaf segmentation and parallel phenotyping for the analysis of gene networks in plants*. In Proc. of European Signal Processing Conference (EUSIPCO 2013), Marrakech, Maroc, September 2013.

[Jeong 07] Y. Jeong & R. J. Radke. *Reslicing axially sampled 3D shapes using elliptic Fourier descriptors*. Medical Image Analysis, vol. 11, no. 2, pages 197–206, Apr 2007.

[Jobson 96] D. J. Jobson, Z. Rahman & G. A. Woodell. *Retinex Image Processing: Improved Fidelity for Direct Visual Observation*. In Proceedings

of the IS&T Fourth Color Imaging Conference: Color Science, Systems, and Applications, 1996.

[Jobson 97] D. J. Jobson, Z. Rahman & G. A. Woodell. *A Multi-Scale Retinex For Bridging the Gap Between Color Images and the Human Observation of Scenes*. IEEE Transactions on Image Processing: Special Issue on Color Processing, 1997.

[Jovanov 11] L. Jovanov, A. Pižurica & W. Philips. *Denoising Algorithm for the 3D Depth Map Sequences Based on Multihypothesis Motion Estimation*. EURASIP Journal on Advances in Signal Processing, 2011.

[Jovanov 13] L. Jovanov, H. Luong & W. Philips. *Depth video enhancement for 3D displays*. In SID Mid-Europe Chapter Spring Meeting 2013, Abstracts, pages 50–52. Society for Information Display, 2013.

[Kagalwala 03] F. Kagalwala & T. Kanade. *Reconstructing specimens using DIC microscope images*. IEEE Transactions on Systems Man and Cybernetics Part B-Cybernetics, vol. 33, no. 5, pages 728–737, 2003.

[Kalman 60] R. Kalman. *A new approach to linear filtering an prediction problems*. Transactions of the ASME-Journal of Basic Engineering, vol. 82, pages 35–45, 1960.

[Kam 98] Zvi Kam. *Microscopic differential interference contrast image processing by line integration (LID) and deconvolution*. Bioimaging, vol. 6, no. 4, pages 166–176, 1998.

[Kamentsky 11] L. Kamentsky, T. R. Jones, A. Fraser, M. A. Bray, D. J. Logan, K. L. Madden, V. Ljosa, C. Rueden, K. W. Eliceiri & A. E. Carpenter. *Improved structure, function and compatibility for CellProfiler: modular high-throughput image analysis software*. Bioinformatics, vol. 27, no. 8, pages 1179–1180, 2011.

[Kapur 85] JN Kapur, Prasanna K Sahoo & AKC Wong. *A new method for gray-level picture thresholding using the entropy of the histogram*. Computer vision, graphics, and image processing, vol. 29, no. 3, pages 273–285, 1985.

[Kass 87] M. Kass, A. Witkin & D. Terzopoulos. *Snakes - Active Contour Models*. International Journal of Computer Vision, vol. 1, no. 4, pages 321–331, 1987.

[Kass 88] M. Kass, A Witkin & D. Terzopoulos. *Snakes: active contour Models*. International journal of computer vision, pages 321–331, 1988.

[Kassim 99] Ashraf A Kassim, T Tan & KH Tan. *A comparative study of efficient generalised Hough transform techniques*. Image and Vision Computing, vol. 17, no. 10, pages 737–748, 1999.

[Kauffman 09] Joost Kauffman. *Automated radiographic assessment of hhand in rheumatoid arthritis*. PhD thesis, Universiteit Twente, 2009.

[Keller 08] P. J. Keller, A. D. Schmidt, J. Wittbrodt & E. H. K. Stelzer. *Reconstruction of Zebrafish Early Embryonic Development by Scanned Light Sheet Microscopy*. Science, vol. 322, no. 5904, pages 1065–1069, Nov 14 2008.

[Keller 10] P. J. Keller, A. D. Schmidt, A. Santella, K. Khairy, Z. R. Bao, J. Wittbrodt & E. H. K. Stelzer. *Fast, high-contrast imaging of animal development with scanned light sheet-based structured-illumination microscopy*. Nature Methods, vol. 7, no. 8, pages 637–U55, Aug 2010.

[Keyes 99] L. Keyes & A. Winstanley. *Fourier Descriptors as A Genereal Classification Tool for Topographic Shapes*. Proceedings Irish Machine Vision and Image Processing Conference, pages 193–203, 1999.

[Kim 00] W. Y. Kim & Y. S. Kim. *A region-based shape descriptor using Zernike moments*. Signal Processing-Image Communication, vol. 16, no. 1-2, pages 95–102, Sep 2000.

[Kim 09] Kee K. Kim, Robert S. Adelstein & Sachiyo Kawamoto. *Identification of Neuronal Nuclei (NeuN) as Fox-3, a New Member of the Fox-1 Gene Family of Splicing Factors*. Journal of Biological Chemistry, vol. 284, no. 45, pages 31052–31061, 2009.

[Kim 11] Il-Han Kim, Y.-C.M. Chen, D.L. Spector, R. Eils & K. Rohr. *Nonrigid Registration of 2-D and 3-D Dynamic Cell Nuclei Images for Improved Classification of Subcellular Particle Motion*. Image Processing, IEEE Transactions on, vol. 20, no. 4, pages 1011–1022, 2011.

[Kindratenko 03] V. V. Kindratenko. *On using functions to describe the shape*. Journal of Mathematical Imaging and Vision, vol. 18, no. 3, pages 225–245, May 2003.

[Klodt 11] M. Klodt & D. Cremers. *A Convex Framework for Image Segmentation with Moment Constraints*. 2011 IEEE International Conference on Computer Vision (ICCV), pages 2236–2243, 2011.

[Klodt 13] Maria Klodt, Frank Steinbrücker & Daniel Cremers. *Moment Constraints in Convex Optimization for Segmentation and Tracking*. In Advanced Topics in Computer Vision, pages 215–242. Springer, 2013.

[Knott 21] L.E. Knott. Apparatus company: Scientific instruments, catalogue 26. 1921.

[Kocsis 91] Eva Kocsis, Benes L Trus, Clifford J Steer, Margaret E Bisher & Alasdair C Steven. *Image averaging of flexible fibrous macromolecules: the clathrin triskelion has an elastic proximal segment*. Journal of structural biology, vol. 107, no. 1, pages 6–14, 1991.

[Koenderink 94] Jan J. Koenderink & Andrea J. van Doorn. *Two-plus-one-dimensional differential geometry*. Pattern Recognition Letters, vol. 15, no. 5, pages 439 – 443, 1994.

[Koopman 06] W. J. H. Koopman, S. Verkaart, S. E. van Emst-de Vries, S. Grefte, J. A. M. Smeitink & P. H. G. M. Willems. *Simultaneous quantification of oxidative stress and cell spreading using 5-(and-6)-chloromethyl-2 ',7 '-dichlorofluorescein*. Cytometry Part A, vol. 69A, no. 12, pages 1184–1192, 2006.

[Krieger 08] J. D. Krieger. *Extended Eigenshape v3.0. Morpho-tools*. Rapport technique, University of Colorado, 2008.

[Krishnan 07] D. Krishnan, P. Lin & A. M. Yip. *A primal-dual active-set method for non-negativity constrained total variation deblurring problems*. IEEE Transactions on Image Processing, vol. 16, no. 11, pages 2766–2777, Nov 2007.

[Kroon 09] D. J. Kroon & C. H. Slump. *MRI Modalitiy Transformation in Demon Registration*. 2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro, Vols 1 and 2, pages 963–966 1405, 2009.

[Kuhn 55] Harold W Kuhn. *The Hungarian method for the assignment problem*. Naval research logistics quarterly, vol. 2, no. 1-2, pages 83–97, 1955.

[Kwok 09] Show Ming Kwok. *In vivo visualization of CaMKII activity in ocular dominance plasticity*. PhD thesis, Massachusetts Institute of Technology (MIT), 2009.

[Lamprecht 07] M. R. Lamprecht, D. M. Sabatini & A. E. Carpenter. *CellProfiler(TM): free, versatile software for automated biological image analysis*. Biotechniques, vol. 42, no. 1, pages 71–75, Jan 2007.

[Leavers 00] V.F. Leavers. *Use of the two-dimensional Radon transform to generate a taxonomy of shape for the characterization of abrasive powder particles*. Pattern Analysis and Machine Intelligence, IEEE Transactions on, vol. 22, no. 12, pages 1411–1423, 2000.

[Leister 99] D. Leister, C. Varotto, P. Pesaresi, A. Niwergall & F. Salamini. *Large-scale evaluation of plant growth in Arabidopsis thaliana by non-invasive image analysis*. Plant Physiology and Biochemistry, vol. 37, no. 9, pages 671–678, 1999.

[Lemaire 06] Patrick Lemaire. *How Many Ways to Make a Chordate?* Science, vol. 312, no. 5777, pages 1145–1146, 2006.

[Lempitsky 07] Victor Lempitsky, Carsten Rother & Andrew Blake. *Logcut-efficient graph cut optimization for markov random fields*. In Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on, pages 1–8. IEEE, 2007.

[Lepez 11] T. Lepez, M. Vandewoestyne, S. Hussain, F. Van Nieuwerburgh, K. Poppe, B. Velkeniers, J. M. Kaufman & D. Deforce. *Fetal Microchimeric Cells in Blood of Women with an Autoimmune Thyroid Disease*. Plos One, vol. 6, no. 12, 2011.

[Lewin 10] Sergej Lewin, Xiaoyi Jiang & Achim Clausing. *Perceptually motivated shape evolution with shape-preserving property*. Pattern Recognition Letters, vol. 31, no. 6, pages 447–453, 2010.

[Li 92] YJ Li. *Reforming the theory of invariant moments for pattern-recognition*. Pattern recognition, vol. 25, no. 7, pages 723–730, JUL 1992.

[Li 07a] B Li & S. Acton. *Active Contour External Force Using Vector Field Convolution for Image Segmentation*. IEEE transactions on image processing, vol. 16, pages 2096–2106, 2007.

[Li 07b] Gang Li, Tianming Liu, Ashley Tarokh, Jingxin Nie, Lei Guo, Andrew Mara, Scott Holley & Stephen TC Wong. *3D cell nuclei segmentation based on gradient flow tracking*. BMC Cell Biology, vol. 8, pages 1–10, 2007.

[Li 08] G. Li, T. Liu, J. Nie, L. Guo, J. Chen, J. Zhu, W. Xia, A. Mara, S. Holley & S. T. C. Wong. *Segmentation of touching cell nuclei using gradient flow tracking*. Journal of Microscopy-Oxford, vol. 231, no. 1, pages 47–58, Jul 2008.

[Li 09] Stan Z Li & S Singh. Markov random field modeling in image analysis, volume 3. Springer, 2009.

[Li 13] Fuhai Li, Zheng Yin, Guangxu Jin, Hong Zhao & Stephen T. C. Wong. *Chapter 17: Bioimage Informatics for Systems Pharmacology*. Genome Research, vol. 9, no. 4, page e1003043, 04 2013.

[Lin 03] Gang Lin, Umesh Adiga, Kathy Olson, John F. Guzowski, Carol A. Barnes & Badrinath Roysam. *A hybrid 3D watershed algorithm incorporating gradient cues and object models for automatic segmentation of nuclei in confocal image stacks*. Cytometry Part A, vol. 56A, no. 1, pages 23–36, 2003.

[Lindeberg 98] Tony Lindeberg. *Edge Detection and Ridge Detection with Automatic Scale Selection*. International Journal of Computer Vision, vol. 30, no. 2, pages 117–156, 1998.

[Liu 07] Maofu Liu, Yanxiang He & Bin Ye. *Image zernike moments shape feature evaluation based on image reconstruction*. Geo-spatial Information Science, vol. 10, no. 3, pages 191–195, 2007.

[Liu 12] Xinyang Liu, Kemal Tuncali, William Wells, Paul Morrison & Gary P. Zientara. *Fully Automatic 3D Segmentation of Iceball for Image-Guided Cryoablation*. In 34th Annual International IEEE EMBS Conference. IEEE, 2012.

[Loncaric 98] S. Loncaric. *A survey of shape analysis techniques*. Pattern Recognition, vol. 31, no. 8, pages 983–1001, Aug 1998.

[Louis 05] S. F. Louis, B. J. Vermolen, Y. Garini, I. T. Young, A. Guffei, Z. Lichtensztejn, F. Kuttler, T. C. Y. Chuang, S. Moshir, V. Mougey, A. Y. C. Chuang, P. D. Kerr, T. Fest, P. Boukamp & S. Mai. *c-Myc induces chromosomal rearrangements through telomere and chromosome remodeling in the interphase nucleus*. Proceedings of the National Academy of Sciences of the United States of America, vol. 102, no. 27, pages 9613–9618, Jul 5 2005.

[Lowe 04] David G Lowe. *Distinctive image features from scale-invariant keypoints*. International journal of computer vision, vol. 60, no. 2, pages 91–110, 2004.

[Luong 14] Hiep Luong. *3D Fluorescence Microscopy Deconvolution Using Alternating Direction Optimization*. In Second International Challenge on 3D Deconvolution Microscopy, IEEE International Symposium on Biomedical Imaging, 2014.

[Mahoor 05] M. H. Mahoor & M. Abdel-Mottaleb. *Classification and numbering of teeth in dental bitewing images*. Pattern Recognition, vol. 38, no. 4, pages 577–586, Apr 2005.

[Malladi 95] R. Malladi, J. A. Sethian & B. C. Vemuri. *Shape Modeling with Front Propagation - a Level Set Approach*. IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 17, no. 2, pages 158–175, 1995.

[Malladi 96] R. Malladi & J. A. Sethian. *Level set and fast marching methods in image processing and computer vision*. International Conference on Image Processing, Proceedings - Vol I, pages 489–492, 1996.

[Mao 10] H. Mao, H. Liu & P. Shi. *A Convex Neighbor-Constrained Active Contour Model For Image Segmentation*. In IEEE International Conference on Image Processing, 2010.

[Marcuzzo 09] M. Marcuzzo, P. Quelhas, A. Campilho, A. M. Mendonca & A. Campilho. *Automated Arabidopsis plant root cell segmentation based on SVM classification and region merging*. Computers in Biology and Medicine, vol. 39, no. 9, pages 785–793, 2009.

[Mat 14] *Correcting Nonuniform Illumination*. Rapport technique, Mathworks, 2014.

[Mattes 06] J. Mattes, J. Nawroth, P. Boukamp, R. Eils & K. M. Greulich-Bode. *Analyzing motion and deformation of the cell nucleus for studying colocalizations of nuclear structures*. 2006 3rd IEEE International Symposium on Biomedical Imaging: Macro to Nano, Vols 1-3, pages 1044–1047, 2006.

[Matula 06] P. Matula, P. Matula, M. Kozubek & V. Dvorak. *Fast point-based 3-D alignment of live cells*. IEEE Transactions on Image Processing, vol. 15, no. 8, pages 2388–2396, Aug 2006.

[McLellan 98] T. McLellan & J. A. Endler. *The relative success of some methods for measuring and describing the shape of complex objects*. Systematic Biology, vol. 47, no. 2, pages 264–281, Jun 1998.

[Megason 07] Sean G. Megason & Scott E. Fraser. *Imaging in Systems Biology*, September 2007.

[Mei 08] Ye Mei & Dimitrios Androutsos. *Affine invariant shape descriptors: the ICA-Fourier descriptor and the PCA-Fourier descriptor*. In Pattern Recognition, 2008. ICPR 2008. 19th International Conference on, pages 1–4. IEEE, 2008.

[Mei 09] Ye Mei & Dimitrios Androutsos. *Robust affine invariant shape image retrieval using the ICA Zernike moment shape descriptor*. In Image Processing (ICIP), 2009 16th IEEE International Conference on, pages 1065–1068. IEEE, 2009.

[Meijering 06] E. Meijering, I. Smal & G. Danuser. *Tracking in molecular bioimaging*. IEEE Signal Processing Magazine, vol. 23, no. 3, pages 46–53, May 2006.

[Meijering 10] E. Meijering. *Neuron Tracing in Perspective*. Cytometry Part A, vol. 77A, no. 7, pages 693–704, 2010.

[Meijering 12] E. Meijering. *Cell Segmentation: 50 Years Down the Road*. IEEE Signal Processing Magazine, vol. 29, no. 5, pages 140–145, 2012.

[Merchant 02] F. A. Merchant & K. R. Castleman. *Strategies for automated fetal cell screening*. Human Reproduction Update, vol. 8, no. 6, pages 509–521, 2002.

[Merks 11] R. M. H. Merks, M. Guravage, D. Inze & G. T. S. Beemster. *VirtualLeaf: An Open-Source Framework for Cell-Based Modeling of Plant Tissue Growth and Development*. Plant Physiology, vol. 155, no. 2, pages 656–666, 2011.

[Meyer 09] M. G. Meyer, M. Fauver, J. R. Rahn, T. Neumann, F. W. Patten, E. J. Seibel & A. C. Nelson. *Automated cell analysis in 2D and 3D: A comparative study*. Pattern Recognition, vol. 42, no. 1, pages 141–146, Jan 2009.

[Meyer 11] T. Meyer, N. Bergner, C. Bielecki, C. Krafft, D. Akimov, B.F. Romeike, R. Reichart, R. Kalff, B. Dietzek & J. Popp. *Nonlinear microscopy, infrared, and Raman microspectroscopy for brain tumor analysis*. Journal of Biomedical Optics, vol. 16, page 0021113, 2011.

[Miao 12] Q. Miao, A.P. Reeves, F.W. Patten & E.J. Seibel. *Multimodal 3D Imaging of Cells and Tissue, Bridging the Gap Between Clinical and Research Microscopy*. Annals of Biomedical Engineering, vol. 40, pages 263–276, 2012.

[Milios 00] E. Milios & E. G. M. Petrakis. *Shape retrieval based on dynamic programming*. IEEE Transactions on Image Processing, vol. 9, no. 1, pages 141–147, Jan 2000.

[Mille 09] J. Mille. *Narrow band region-based active contours and surfaces for 2D and 3D segmentation*. Computer Vision and Image Understanding, vol. 113, no. 9, pages 946–965, Sep 2009.

[Mokhtarian 97] Farzin Mokhtarian, Sadegh Abbasi & Josef Kittler. *Efficient and robust retrieval by shape content through curvature scale space*. Series on Software Engineering and Knowledge Engineering, vol. 8, pages 51–58, 1997.

[Mokhtarian 11] Farzin Mokhtarian & Miroslaw Bober. Curvature scale space representation: theory, applications, and mpeg-7 standardization. Springer Publishing Company, Incorporated, 2011.

[Molenaar 03] C. Molenaar, K. Wiesmeijer, N. P. Verwoerd, S. Khazen, R. Eils, H. J. Tanke & R. W. Dirks. *Visualizing telomere dynamics in living mammalian cells using PNA probes*. EMBO Journal, vol. 22, no. 24, pages 6631–6641, 2003.

[Mori 05] Greg Mori, Serge Belongie & Jitendra Malik. *Efficient shape matching using shape contexts*. Pattern Analysis and Machine Intelligence, IEEE Transactions on, vol. 27, no. 11, pages 1832–1837, 2005.

[Mortensen 98] Eric N Mortensen & William A Barrett. *Interactive segmentation with intelligent scissors*. Graphical models and image processing, vol. 60, no. 5, pages 349–384, 1998.

[Muchir 03] A. Muchir, B. G. van Engelen, M. Lammens, J. M. Mislow, E. McNally, K. Schwartz & G. Bonne. *Nuclear envelope alterations in fibroblasts from LGMD1B patients carrying nonsense Y259X heterozygous or homozygous mutation in lamin A/C gene*. Exp Cell Res, vol. 291, no. 2, pages 352–62, 2003.

[Mukundan 00] R. Mukundan, S. H. Ong & P. A. Lee. *Discrete Orthogonal Moment Features Using Chebyshev Polynomials*. In New Zealand: International Conference on Image and Vision Computing - IVCNZ'00, 2000.

[Mullen 06] Jack L Mullen, Cynthia Weinig & Roger P Hangarter. *Shade avoidance and the regulation of leaf inclination in Arabidopsis*. Plant, Cell & Environment, vol. 29, no. 6, pages 1099–1106, 2006.

[Mumford 89] D. Mumford & J. Shah. *Optimal approximations by piecewise smooth functions and associated variational problems*. Communications on Pure and Applied Mathematics, vol. 42, no. 5, pages 577–685, 1989.

[Murtagh 03] Fionn Murtagh, J.L. Starck, E.J. Candes & D.L. Donoho. *Gray and color image contrast enhancement by the curvelet transform*. IEEE Transactions on Image Processing, vol. 12(6), pages 706–717, 2003.

[Myronenko 10] A. Myronenko & X. B. Song. *Intensity-Based Image Registration by Minimizing Residual Complexity*. IEEE Transactions on Medical Imaging, vol. 29, no. 11, pages 1882–1891, 2010.

[Nakhmani 12] A. Nakhmani & A. Tannenbaum. *Self-Crossing Detection and Location for Parametric Active Contours*. Image Processing, IEEE Transactions on, vol. 21, no. 7, pages 3150–3156, July 2012.

[Nemeth 11] J. Nemeth, Z. Kato & I. Jermyn. *A Multi-Layer 'Gas of Circles' Markov Random Field Model for the Extraction of Overlapping Near-Circular Objects*. In ACIVS, pages 171–182, 2011.

[Neto 06] J. C. Neto, G. E. Meyer, D. D. Jones & A. K. Samal. *Plant species identification using Elliptic Fourier leaf shape analysis*. Computers and Electronics in Agriculture, vol. 50, no. 2, pages 121–134, Feb 2006.

[Nieuwenhuis 13a] Claudia Nieuwenhuis & Daniel Cremers. *Spatially varying color distributions for interactive multilabel segmentation*. Pattern Analysis and Machine Intelligence, IEEE Transactions on, vol. 35, no. 5, pages 1234–1247, 2013.

[Nieuwenhuis 13b] Claudia Nieuwenhuis, Evgeny Strekalovskiy & Daniel Cremers. *Proportion Priors for Image Sequence Segmentation*. In ICCV, pages 2328–2335, 2013.

[Nieuwenhuis 13c] Claudia Nieuwenhuis, Eno Töppe & Daniel Cremers. *A survey and comparison of discrete and continuous multi-label optimization approaches for the potts model*. International journal of computer vision, vol. 104, no. 3, pages 223–240, 2013.

[Obara 11] Boguslaw Obara, Michael Veeman, Jae Hyeok Choi, William Smith & B. S. Manjunath. *Segmentation of ascidian notochord cells in DIC timelapse images*. Microscopy Research and Technique, vol. 74, pages 727–734, 2011.

[Ochoa Donoso 10] Daniel Ochoa Donoso, Sidharta Gautama & Wilfried Philips. *Automatic identification of Caenorhabditis elegans in population images by shape energy features*. J Microsc, vol. 238, no. 2, pages 173–84, May 2010.

[Ochoa Donoso 11] Daniel Ochoa Donoso. *Biological image analysis of model organisms*. PhD thesis, Ghent University. Faculty of Engineering and Architecture, 2011.

[Ollion 13] Jean Ollion, Julien Cochennec, François Loll, Christophe Escudé & Thomas Boudier. *TANGO: a generic tool for high-throughput 3D image analysis for studying nuclear organization*. Bioinformatics, vol. 29, no. 14, pages 1840–1841, 2013.

[Onyango 01] C. M. Onyango & J. A. Marchant. *Physics-based colour image segmentation for scenes containing vegetation and soil*. Image and Vision Computing, vol. 19, no. 8, pages 523–538, Jun 1 2001.

[Onyango 03] C. M. Onyango & J. A. Marchant. *Segmentation of row crop plants from weeds using colour and morphology*. Computers and Electronics in Agriculture, vol. 39, no. 3, pages 141–155, Aug 2003.

[Orjuela-Vargas 13] Sergio Alejandro Orjuela-Vargas. *Texture analysis for the evaluation of appearance changes in textile surfaces*. PhD thesis, Ghent university, 2013.

[Orjuela-Vargas 14a] S. Orjuela-Vargas & W. Philips. *Survey on texture analysis using LBP techniques*. In SPIE Optics, Photonics and Digital Technologies for Multimedia Applications, volume 9138, pages 9138–18, 2014.

[Orjuela-Vargas 14b] Sergio Alejandro Orjuela-Vargas, Juan Pablo Yañez & Wilfried Philips. The geometric local textural patterns (gltp) technique, volume 506 of *Local binary patterns : new variants and new applications*, pages 30–70. Springer, 2014.

[Otsu 76] N. Otsu. *A threshold Selection Method from Gray Level Histograms*. IEEE transactions on Systems, Man and Cybernetics, vol. 9, pages 62–66, 1976.

[Otsu 79] N. Otsu. *Threshold Selection Method from Gray-Level Histograms*. IEEE Transactions on Systems Man and Cybernetics, vol. 9, no. 1, pages 62–66, 1979.

[Padfield 11] Dirk Padfield, Jens Rittscher & Badrinath Roysam. *Coupled minimum-cost flow cell tracking for high-throughput quantitative analysis*. Medical Image Analysis, vol. 15, no. 4, pages 650–668, 2011.

[Papandreou 07] G. Papandreou & P. Maragos. *Multigrid geometric active contour models*. IEEE Transactions on Image Processing, vol. 16, no. 1, pages 229–240, 2007.

[Paragios 02] Nikos Paragios & Rachid Deriche. *Geodesic active regions and level set methods for supervised texture segmentation*. International Journal of Computer Vision, vol. 46, no. 3, pages 223–247, 2002.

[Paramita 12] B. Paramita & P. Anupam. *A new tool for analysis of root growth in the spatio-temporal continuum*. New Phytologist, vol. 195, pages 264–274, 2012.

[Park 02] Hyun Keun Park & Myung Jin Chung. *External force of snake: virtual electric field*. Electronics Letters, vol. 38, no. 24, pages 1500–1502, 2002.

[Petrakis 02] E. G. M. Petrakis, A. Diplaros & E. Milios. *Matching and retrieval of distorted and occluded shapes using dynamic programming*. IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 24, no. 11, pages 1501–1516, Nov 2002.

[Peyre 08] G. Peyre. *A toolbox for fast marching and level set computations*. Rapport technique, 2008.

[Pfeffer 03] W. Pfeffer. The physiology of plants (edn 2), vol 2. Oxford, 1903.

[Philips 93] W. Philips. *A New Fast Algorithm for Moment Computation*. Pattern recognition, vol. 26, no. 11, pages 1619–1621, 1993.

[Philips 95] W. Philips. *Adaptive warped polynomial contour descriptors*. 1995 IEEE International Conference on Systems, Man and Cybernetics, Vols 1-5, pages 4638–4643 4722, 1995.

[Philips 96] W. Philips. *Fast coding of arbitrarily shaped image segments using weakly separable bases*. Optical Engineering, vol. 35, no. 1, pages 177–186, Jan 1996.

[Pižurica 02] A. Pižurica. *Image Denoising Using Wavelets and ModellinContext Modelling*. PhD thesis, Ghent University, 2002.

[Pock 10] T. Pock, D. Cremers, H. Bischof & A. Chambolle. *Global Solutions of Variational Models with Convex Regularization*. Siam Journal on Imaging Sciences, vol. 3, no. 4, pages 1122–1145, 2010.

[Poon 97] C. S. Poon & M. Braun. *Image segmentation by a deformable contour model incorporating region analysis*. Physics in Medicine and Biology, vol. 42, no. 9, pages 1833–1841, Sep 1997.

[Potts 52] Renfrey Burnard Potts. *Some generalized order-disorder transformations*. In Proceedings of the Cambridge Philosophical Society, volume 48, pages 106–109. Cambridge Univ Press, 1952.

[Pound 12] M. P. Pound, A. P. French, D. M. Wells, M. J. Bennett & T. P. Pridmore. *CellSeT: Novel Software to Extract and Analyze Structured Networks of Plant Cells from Confocal Images*. Plant Cell, vol. 24, no. 4, pages 1353–1361, 2012.

[Price 11] C. A. Price, O. Symonova, Y. Mileyko, T. Hilley & J. S. Weitz. *Leaf extraction and analysis framework graphical user interface: segmenting and analyzing the structure of leaf veins and areoles*. Plant Physiol, vol. 155, no. 1, pages 236–45, Jan 2011.

[Pridmore 12] Tony Pridmore. *X-Ray micro-Computed Tomography for Plant Root Phenotyping*. In Phenodays, 2012.

[Purbasha 09] S. Purbasha, E. Bosneaga & A. Manfred. *Plant cell walls throughout evolution: towards a molecular understanding of their design principles*. Journal of Experimental Botany, vol. 60, pages 3615–3635, 2009.

[Quelhas 10a] P. Quelhas, M. Marcuzzo, A. M. Mendonca & A. Campilho. *Cell Nuclei and Cytoplasm Joint Segmentation Using the Sliding Band Filter*. IEEE Transactions on Medical Imaging, vol. 29, no. 8, pages 1463–1473, Aug 2010.

[Quelhas 10b] Pedro Quelhas, Ana Maria Mendonça & Aurélio Campilho. *Optical Flow Based Arabidopsis Thaliana Root Meristem Cell Division Detection*. In Aurélio Campilho & Mohamed Kamel, editeurs, Image Analysis and Recognition, volume 6112 of *Lecture Notes in Computer Science*, pages 217–226. Springer Berlin Heidelberg, 2010.

[Radke 04] R. Radke, D. Freedman, T. Zhang, Y. Jeong & G. Chen. *Deformable model-based segmentation of 3D CT by matching distributions*. Medical Physics, vol. 31, no. 6, pages 1846–1846, Jun 2004.

[Ray 04] N. Ray & S. T. Acton. *Motion gradient vector flow: An external force for tracking rolling leukocytes with shape and size constrained active contours*. IEEE Transactions on Medical Imaging, vol. 23, no. 12, pages 1466–1478, Dec 2004.

[Raz 06] V. Raz, F. Carlotti, B. J. Vermolen, E. van der Poel, W. C. Sloos, S. Knaan-Shanzer, A. A. de Vries, R. C. Hoeben, I. T. Young, H. J. Tanke, Y. Garini & R. W. Dirks. *Changes in lamina structure are followed by spatial reorganization of heterochromatic regions in caspase-8-activated human mesenchymal stem cells*. J Cell Sci, vol. 119, no. Pt 20, pages 4247–56, Oct 15 2006.

[Raz 08] V. Raz, B. J. Vermolen, Y. Garini, J. J. M. Onderwater, M. A. Mommaas-Kienhuis, A. J. Koster, I. T. Young, H. Tanke & R. W. Dirks. *The nuclear lamina promotes telomere aggregation and centromere peripheral localization during senescence of human mesenchymal stem cells*. Journal of Cell Science, vol. 121, no. 24, pages 4018–4028, Dec 15 2008.

[Rembold 06] Martina Rembold, Felix Loosli, Richard J. Adams & Joachim Wittbrodt. *Individual Cell Migration Serves as the Driving Force for Optic Vesicle Evagination*. Science, vol. 313, no. 5790, pages 1130–1134, 2006.

[Ridler 78] TW Ridler & S Calvard. *Picture thresholding using an iterative selection method*. IEEE transactions on Systems, Man and Cybernetics, vol. 8, no. 8, pages 630–632, 1978.

[Rieger 04] B. Rieger, C. Molenaar, R. W. Dirks & L. J. Van Vliet. *Alignment of the cell nucleus from labeled proteins only for 4D in vivo imaging*.

Microscopy Research and Technique, vol. 64, no. 2, pages 142–150, Jun 1 2004.

[Roberts 07] T. Roberts, S. McKenna, N. Wuyts, T. Valentine & A. Bengough. *Performance of Low-Level Motion Estimation Methods for Confocal Microscopy of Plant Cells in vivo*. In Motion and Video Computing, 2007. WMVC '07. IEEE Workshop on, pages 13–13, 2007.

[Rochery 06] Marie Rochery, Ian H. Jermyn & Josiane Zerubia. *Higher Order Active Contours*. Int. J. Comput. Vision, vol. 69, no. 1, pages 27–42, 2006.

[Roels 14] Joris Roels. Beeldoptimalisatie voor 3d elektronenmicroscopie beelden. Master's thesis, Ghent university, 2014.

[Rohde 08] G. K. Rohde, A. J. Ribeiro, K. N. Dahl & R. F. Murphy. *Deformation-based nuclear morphometry: capturing nuclear shape variation in HeLa cells*. Cytometry A, vol. 73, no. 4, pages 341–50, Apr 2008.

[Rooms 05a] F. Rooms, W. Philips & D. Lidke. *Simultaneous degradation estimation and restoration of confocal images and performance evaluation by colocalization analysis*. Journal of Microscopy, vol. 218, no. 1, pages 22–36, April 2005.

[Rooms 05b] Filip Rooms. *Nonlinear methods in image restoration*. PhD thesis, Ghent University. Faculty of Engineering and Architecture, 2005.

[Rubinstein 10] R. Rubinstein, A.M. Bruckstein & M. Elad. *Dictionaries for Sparse Representation Modeling*. Proceedings of the IEEE, vol. 98, no. 6, pages 1045–1057, June 2010.

[Rubner 00] Y. Rubner, C. Tomasi & L. J. Guibas. *The Earth Mover's Distance as a metric for image retrieval*. International Journal of Computer Vision, vol. 40, no. 2, pages 99–121, Nov 2000.

[Ruusuvuori 08] P. Ruusuvuori, A. Lehmussola, J. Selinummi, T. Rajala, H. Huttunen & O. Yli-Harja. *set of synthetic images for validating cell image analysis*. Proc. of the 16th European Signal Processing Conference (EUSIPCO-2008), 2008.

[Ružić 14] Tijana Ružić. *Patch-Based Graphical Models for Image Restoration*. PhD thesis, Ghent university, 2014.

[Saad 03] Y. Saad. Iterative methods for sparse linear systems. SIAM, Philadelphia, 2nd edition, 2003.

[Schmundt 98] D. Schmundt, M. Stitt, B. Jahne & U. Schurr. *Quantitative analysis of the local rates of growth of dicot leaves at a high temporal and spatial resolution, using image sequence analysis*. Plant Journal, vol. 16, no. 4, pages 505–514, 1998.

[Schoenemann 07] T. Schoenemann & D. Cremers. *Globally optimal image segmentation with an elastic shape prior*. 2007 IEEE 11th International Conference on Computer Vision, Vols 1-6, pages 1132–1137, 2007.

[Schrama 06] David Schrama, Ralph A. Reisfeld & Jurgen C. Becker. *Antibody targeted drugs as cancer therapeutics*. Nat Rev Drug Discov, vol. 5, no. 2, pages 147–159, February 2006.

[Schulz 12] Hannes Schulz, Johannes A Postma, Dagmar van Dusschoten, Hanno Scharr & Sven Behnke. *3D Reconstruction of Plant Roots from MRI Images.* In VISAPP (2), pages 24–33, 2012.

[Selinummi 05] J. Selinummi, J. Seppala, O. Yli-Harja & J. A. Puhakka. *Software for quantification of labeled bacteria from digital microscope images by automated image analysis*. Biotechniques, vol. 39, no. 6, pages 859–863, Dec 2005.

[Sener 06] S. Sener & M. Unel. *A new affine invariant curve normalization technique using independent component analysis*. 18th International Conference on Pattern Recognition, Vol 2, Proceedings, pages 48–51 1276, 2006.

[Sethian 99] James Albert Sethian. Level set methods and fast marching methods : evolving interfaces in computational geometry, fluid mechanics, computer vision, and materials science. Cambridge monographs on applied and computational mathematics 3. Cambridge University Press, Cambridge, U.K. ; New York, 2nd edition, 1999.

[Sethian 02] J. A. Sethian. *A review of level set and fast marching methods for image processing*. Modern Methods in Scientific Computing and Applications, vol. 75, pages 365–396, 2002.

[Setzer 09] Simon Setzer. *Split Bregman Algorithm, Douglas-Rachford Splitting and Frame Shrinkage*. In Xue-Cheng Tai, Knut Marken, Marius Lysaker & Knut-Andreas Lie, editeurs, Scale Space and Variational Methods in Computer Vision, volume 5567 of *Lecture Notes in Computer Science*, pages 464–476. Springer Berlin Heidelberg, 2009.

[Sezgin 04] M. Sezgin & B. Sankur. *Survey over image thresholding techniques and quantitative performance evaluation*. Journal of Electronic Imaging, vol. 13, no. 1, pages 146–168, Jan 2004.

[Shamir 10] L. Shamir, J. D. Delaney, N. Orlov, D. M. Eckley & I. G. Goldberg. *Pattern Recognition Software and Techniques for Biological Image Analysis*. PLOS Computational Biology, vol. 6, no. 11, 2010.

[Shang 11] Yanfeng Shang, R. Deklerck, E. Nyssen, A. Markova, J. de Mey, Xin Yang & Kun Sun. *Vascular Active Contour for Vessel Tree Segmentation*. Biomedical Engineering, IEEE Transactions on, vol. 58, no. 4, pages 1023 –1032, april 2011.

[Shen 96] Jun Shen & Danfei Shen. *Orthogonal Legendre moments and their calculation*. In Pattern Recognition, 1996., Proceedings of the 13th International Conference on, volume 2, pages 241–245 vol.2, 1996.

[Shimizu 95] H. Shimizu & R. D. Heins. *Computer-vision-based system for plant growth analysis*. Trans ASABE, vol. 38, pages 959–964, 1995.

[Shu 13] J. Shu, H. Fu, G. Qiu, P. Kaye & M. Ilyas. *Segmenting Overlapping Cell Nuclei In Digital Histopathologyb Images*. In 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, pages 5445–5448, 2013.

[Sirault 12] Xavier Sirault. *3D plant analysis over time: understanding plant architecture, growth and function using imaging technologies*. In Phenodays, 2012.

[Skloot 10] Rebecca Skloot. The immortal life of henrietta lacks. Broadway Paperbacks, New York, 2010.

[Sofka 06] Michal Sofka & Charles V Stewart. *Retinal vessel centerline extraction using multiscale matched filters, confidence and edge measures*. Medical Imaging, IEEE Transactions on, vol. 25, no. 12, pages 1531–1546, 2006.

[Staib 92] L. H. Staib & J. S. Duncan. *Boundary Finding with Parametrically Deformable Models*. IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 14, no. 11, pages 1061–1075, Nov 1992.

[Steger 98] C. Steger. *An Unbiased Detector of Curvilinear Structures*. IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 20, no. 2, pages 113–125, 1998.

[Stelzer 98] E. H. K. Stelzer. *Contrast, resolution, pixelation, dynamic range and signal-to-noise ratio: fundamental limits to resolution in fluorescence light microscopy*. Journal of Microscopy-Oxford, vol. 189, pages 15–24, Jan 1998.

[Stuhmer 13] Jan Stuhmer, Peter Schroder & Daniel Cremers. *Tree Shape Priors with Connectivity Constraints Using Convex Relaxation on General Graphs*. In The IEEE International Conference on Computer Vision (ICCV), December 2013.

[Suk 11] Tomas Suk & Jan Flusser. *Affine moment invariants generated by graph method*. pattern recognition, vol. 44, no. 9, SI, pages 2047–2056, SEP 2011.

[Sum 07] KW Sum & Paul Cheung. *Boundary vector field for parametric active contours*. Pattern Recognition, vol. 40, no. 6, pages 1635–1645, 2007.

[Sun 09] C. M. Sun, P. Vallotton, D. D. Wang, J. Lopez, Y. Ng & D. James. *Membrane boundary extraction using circular multiple paths*. Pattern Recognition, vol. 42, no. 4, pages 523–530, 2009.

[Sun 12] Da-Wen Sun, Corrado Costa & Paolo Menesatti. *Advantages of Using Quantitative Shape Descriptors in Protocols for Plant Cultivar and Postharvest Product Quality Assessment*. Food and Bioprocess Technology, vol. 5, no. 1, pages 1–2, 2012.

[Sundaramoorthi 09] G. Sundaramoorthi, A. Yezzi, A. C. Mennucci & G. Sapiro. *New Possibilities with Sobolev Active Contours*. International Journal of Computer Vision, vol. 84, no. 2, pages 113–129, 2009.

[Tabbone 06] S. Tabbone, L. Wendling & J.-P. Salmon. *A new shape descriptor defined on the Radon transform*. Computer Vision and Image Understanding, vol. 102, no. 1, pages 42 – 51, 2006.

[Tabus 10] Jenni Hukkanen; Andrea Hategan; Edmond Sabo; Ioan Tabus. *Segmentation of Cell Nuclei from Histological Images by Ellipse Fitting*. 2010.

[Tahmasbi 11] Amir Tahmasbi, Fatemeh Saki & Shahriar B Shokouhi. *Classification of benign and malignant masses based on Zernike moments*. Computers in biology and medicine, vol. 41, no. 8, pages 726–735, 2011.

[Takahashi 90] K. Takahashi, Y. Sawasaki, J. Hata, K. Mukai & T. Goto. *Spontaneous transformation and immortalization of human endothelial cells*. In Vitro Cell Dev Biol, vol. 26, no. 3 Pt 1, pages 265–74, 1990.

[Tang 08] J. Tang. *A multi-direction GVF snake for the segmentation of skin cancer images*. Pattern Recognition, 2008.

[Teague 80] Michael Reed Teague. *Image analysis via the general theory of moments*. J. Opt. Soc. Am., vol. 70, no. 8, pages 920–930, Aug 1980.

[Tessens 10] Linda Tessens. *Information selection and fusion in vision systems*. PhD thesis, Ghent University, 2010.

[Tiehua Du 09] Martin Wasser Tiehua Du. *3D Image Stack Reconstruction in Live Cell Microscopy of Drosophila Muscles and its Validation*. Cytometry A, vol. 75, pages 329–343, 2009.

[Tkaczyk 10] Thomas Tkaczyk. Field guide to microscopy. SPIE, 2010.

[Tosic 11] I. Tosic & P. Frossard. *Dictionary Learning*. Signal Processing Magazine, IEEE, vol. 28, no. 2, pages 27–38, March 2011.

[Trier 96] O. D. Trier, A. K. Jain & T. Taxt. *Feature extraction methods for character recognition - A survey*. Pattern Recognition, vol. 29, no. 4, pages 641–662, Apr 1996.

[Trimeche 00] M. Trimeche, F. A. Cheikh & M. Gabbouj. *Similarity retrieval of occluded shapes using wavelet-based shape features*. Internet Multimedia Management Systems, vol. 4210, pages 281–289 436, 2000.

[Tsai 85] Wen-Hsiang Tsai. *Moment-preserving thresolding: A new approach*. Computer Vision, Graphics, and Image Processing, vol. 29, no. 3, pages 377–393, 1985.

[Tsechpenakis 04] G. Tsechpenakis, K. Rapantzikos, N. Tsapatsoulis & S. Kollias. *A snake model for object tracking in natural sequences*. Signal Processing-Image Communication, vol. 19, no. 3, pages 219–238, Mar 2004.

[Tseng 04] Y. Tseng, J. S. H. Lee, T. P. Kole, I. Jiang & D. Wirtz. *Micro-organization and visco-elasticity of the interphase nucleus revealed by particle nanotracking*. Journal of Cell Science, vol. 117, no. 10, pages 2159–2167, 2004.

[Tsui 10] Po-Hsiang Tsui, Yin-Yin Liao, Chien-Cheng Chang, Wen-Hung Kuo, King-Jen Chang & Chih-Kuang Yeh. *Classification of Benign and Malignant Breast Tumors by 2-D Analysis Based on Contour Description and Scatterer Characterization*. Medical Imaging, IEEE Transactions on, vol. 29, no. 2, pages 513–522, 2010.

[UN 12] UN. *World Population Prospects: The 2012 Revision, Methodology of the United Nations Population Estimates and Projections.* Rapport technique, UN Department of Economic and Social Affairs, Population Division, 2012.

[Van de Wouwer 99] Gert Van de Wouwer, Paul Scheunders & Dirk Van Dyck. *Statistical texture characterization from discrete wavelet representations*. Image Processing, IEEE Transactions on, vol. 8, no. 4, pages 592–598, 1999.

[Van Otterloo 88] P.J. Van Otterloo. *A contour-oriented approach to digital shape analysis*. PhD thesis, Technical University of Delft, 1988.

[Vandewoestyne 09a] M. Vandewoestyne, D. Van Hoofstat, F. Van Nieuwerburgh & D. Deforce. *Automatic detection of spermatozoa for laser capture microdissection*. International Journal of Legal Medicine, vol. 123, no. 2, pages 169–175, 2009.

[Vandewoestyne 09b] M. Vandewoestyne, D. Van Hoofstat, F. Van Nieuwerburgh & D. Deforce. *Suspension fluorescence in situ hybridization (S-FISH) combined with automatic detection and laser microdissection for STR profiling of male cells in male/female mixtures*. International Journal of Legal Medicine, vol. 123, no. 5, pages 441–447, 2009.

[Vicente 08] Sara Vicente, Vladimir Kolmogorov & Carsten Rother. *Graph cut based image segmentation with connectivity priors*. In Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on, pages 1–8. IEEE, 2008.

[Vigouroux 03] B. Vigouroux & A. Clement. *Unsupervised segmentation of scenes containing vegetation (Forsythia) and soil by hierarchical analysis of bi-dimensional histograms*. Pattern Recognition Letters, vol. 24, no. 12, pages 1951–1957, Aug 2003.

[Vincent 91] L. Vincent & P. Soille. *Watersheds in digital spaces: An efficient algorithm based on immersion simulations*. IEEE PAMI, 1991, vol. 13, no. 6, pages 583–598, 1991.

[Vincent 92] L. Vincent & B. Masters. *Morphological image processing and network analysis of cornea endothelial cell images*. Proc. SPIE, vol. 1769, pages 212–226, 1992.

[Vonesch 06] C. Vonesch, F. Aguet, J-L. Vonesch & V. Unser. *The colored revolution of bioimaging*. IEEE signal processing magazine, pages 20–31, 2006.

[Vonesch 13] Cedric Vonesch, Stamatis Lefkimmiatis, Laure Blanc-Feraud, Michael Unser, Rainer Heintzmann & Arne Seitz. *3D Deconvolution Microscopy Challenge*. In 2013 IEEE 10th International Symposium On Biomedical Imaging (ISBI), 2013.

[Walter 07] A. Walter, H. Scharr, F. Gilmer, R. Zierer, K. A. Nagel, M. Ernst, A. Wiese, O. Virnich, M. Christ, B. Uhlig, S. Junger & U. Schurr. *Dynamics of seedling growth acclimation towards altered light conditions can be quantified via GROWSCREEN: a setup and procedure designed for rapid optical phenotyping of different plant species*. New Phytol, vol. 174, no. 2, pages 447–55, 2007.

[Walter 09] A. Walter, M. Jansen, F. Gilmer, B. Biskup, K. A. Nagel, U. Rascher, A. Fischbach, S. Briem, G. Dreissen, S. Tittmann, S. Braun, I. De Jaeger, M. Metzlaff, U. Schurr & H. Scharr. *Simultaneous phenotyping of leaf growth and chlorophyll fluorescence via GROWSCREEN FLUORO allows detection of stress tolerance in Arabidopsis thaliana and other rosette plants*. Functional Plant Biology, vol. 36, no. 10-11, pages 902–914, 2009.

[Wang 03] Wenjie Wang, Yangang Zu, Fengjian Yang, Huimei Wang & Fei Wang. *Photosynthetic ecophysiological study on the growth of korean pine (Pinus koraiensis) afforested by the edge-effect belt method*. Acta Ecologica Sinica, vol. 23, no. 11, pages 2318–2326, 2003.

[Wang 08] Y. Q. Wang & Y. D. Jia. *External Force for Active Contours: Gradient Vector Convolution*. Pricai 2008: Trends in Artificial Intelligence, vol. 5351, pages 466–472 1130, 2008.

[Wang 09] L. Y. Wang, I. V. Uilecan, A. H. Assadi, C. A. Kozmik & E. P. Spalding. *HYPOTrace: Image Analysis Software for Measuring Hypocotyl Growth and Shape Demonstrated on Arabidopsis Seedlings Undergoing Photomorphogenesis*. Plant Physiology, vol. 149, no. 4, pages 1632–1637, 2009.

[Weckler 04] Paul R Weckler, Niels O Maness & Marvin L Stone. *Estimating Water Stress in Plants Using Hyperspectral Sensing*. system, vol. 1, page 4, 2004.

[Wei 09] C. H. Wei, Y. Li, W. Y. Chau & C. T. Li. *Trademark image retrieval using synthetic features for describing global shape and interior structure*. Pattern Recognition, vol. 42, no. 3, pages 386–394, Mar 2009.

[Weiner 10] Louis M. Weiner, Rishi Surana & Shangzi Wang. *Monoclonal antibodies: versatile platforms for cancer immunotherapy*. Nat Rev Immunol, vol. 10, no. 5, pages 317–327, May 2010.

[Weyland 90] NBJ Weyland & R Prasad. *Characterisation of line drawings using generalised Fourier descriptors*. Electronics Letters, vol. 26, no. 21, pages 1794–1795, 1990.

[Wilson 06] C. A. Wilson & J. A. Theriot. *A correlation-based approach to calculate rotation and translation of moving cells*. IEEE Transactions on Image Processing, vol. 15, no. 7, pages 1939–1951, Jul 2006.

[Wilt 09] B.A. Wilt, L.D. Burns, E.T. Ho, K.K. Ghosh, E. A. Mukamel & M. J. Schnitzer. *Advances in Light Microscopy for Neuroscience*. Annu Rev Neurosci., vol. 32, pages 435–506, 2009.

[Wirtz 09] D. Wirtz. *Particle-Tracking Microrheology of Living Cells: Principles and Applications*. Annual Review of Biophysics, vol. 38, pages 301–326, 2009.

[Wood 96] J. Wood. *Invariant Pattern recognition: a review*. Pattern recognition, vol. 29, no. 1, pages 1–17, 1996.

[Xia 06] M. Xia & B. Liu. *Aligning curves under projective transform and its application to image registration*. 2006 IEEE International Conference on Image Processing, ICIP 2006, Proceedings, pages 349–352 3336, 2006.

[Xia 12] Wei Xia, Zheng Song, Jiashi Feng, Loong-Fah Cheong & Shuicheng Yan. *Segmentation over detection by coupled global and local sparse representations*. In Computer Vision–ECCV 2012, pages 662–675. Springer, 2012.

[Xu 98] C. Y. Xu & J. L. Prince. *Snakes, shapes, and gradient vector flow*. IEEE Transactions on Image Processing, vol. 7, no. 3, pages 359–369, Mar 1998.

[Xu 00] C. Y. Xu, A. Yezzi & J. L. Prince. *On the relationship between parametric and geometric active contours*. Conference Record of the Thirty-Fourth Asilomar Conference on Signals, Systems & Computers, pages 483–489 1882, 2000.

[Xu 08] Dong Xu & Hua Li. *Geometric moment invariants*. Pattern recognition, vol. 41, no. 1, pages 240–249, 2008.

[Yadav 07] R. B. Yadav, N. K. Nishchal, A. K. Gupta & V. K. Rastogi. *Retrieval and classification of shape-based objects using Fourier, generic Fourier, and wavelet-Fourier descriptors technique: A comparative study*. Optics and Lasers in Engineering, vol. 45, no. 6, pages 695–708, Jun 2007.

[Yan 08] P. K. Yan, X. B. Zhou, M. Shah & S. T. C. Wong. *Automatic segmentation of high-throughput RNAi fluorescent cellular images*. IEEE Transactions on Information Technology in Biomedicine, vol. 12, no. 1, pages 109–117, Jan 2008.

[YanFeng 11] Shang YanFeng. *Deterministic and probabilistic active contour models for medical image segmentation*. PhD thesis, Vrije Universiteit Brussel, 2011.

[Yang 08] S. W. Yang, D. Kohler, K. Teller, T. Cremer, P. Le Baccon, E. Heard, R. Eils & K. Rohr. *Nonrigid registration of 3-D multichannel microscopy images of cell nuclei*. IEEE Transactions on Image Processing, vol. 17, no. 4, pages 493–499, Apr 2008.

[Yang 10a] A. Y. Yang, S. S. Sastry, A. Ganesh & Ma Yi. *Fast l1-minimization algorithms and an application in robust face recognition: A review*. In Image Processing (ICIP), 2010 17th IEEE International Conference on, pages 1849–1852, 26-29 Sept. 2010 2010.

[Yang 10b] Jianchao Yang, John Wright, Thomas S Huang & Yi Ma. *Image super-resolution via sparse representation*. Image Processing, IEEE Transactions on, vol. 19, no. 11, pages 2861–2873, 2010.

[Yap 08] Choon Kong Yap & Hwee Kuan Lee. *Identification of Cell Nucleus Using a Mumford-Shah Ellipse Detector*. In Proceedings of the 4th International Symposium on Advances in Visual Computing, ISVC '08, pages 582–593, Berlin, Heidelberg, 2008. Springer-Verlag.

[Yip 09] A. M. Yip, D. Krishnan & Q. V. Pham. *A primal-dual active-set algorithm for bilaterally constrained total variation deblurring and piecewise constant Mumford-Shah segmentation problems*. Advances in Computational Mathematics, vol. 31, no. 1-3, pages 237–266, Oct 2009.

[Yoshioka 04] Yosuke Yoshioka, Hiroyoshi Iwata, RYO Ohsawa & Seishi Ninomiya. *Analysis of petal shape variation of Primula sieboldii by elliptic Fourier descriptors and principal component analysis*. Annals of Botany, vol. 94, no. 5, pages 657–664, 2004.

[Younes 99] L. Younes. *Optimal matching between shapes via elastic deformations*. Image and Vision Computing, vol. 17, no. 5-6, pages 381–389, Apr 1999.

[Yu 05] Z. H. Yu & G. Herman. *On the earth mover's distance as a histogram similarity metric for image retrieval*. 2005 IEEE International Conference on Multimedia and Expo (ICME), Vols 1 and 2, pages 686–689 1598, 2005.

[Yu 09] D. G. Yu, T. D. Pham, X. B. Zhou & S. T. C. Wong. *Recognition and analysis of cell nuclear phases for high-content screening based on morphological features*. Pattern Recognition, vol. 42, no. 4, pages 498–508, Apr 2009.

[Zernike 34] F. Zernike. *Beugungstheorie des schneidenver-fahrens und seiner verbesserten form, der phasenkontrastmethode*. Physica, vol. 1, no. 7, pages 689 – 704, 1934.

[Zeyde 12] Roman Zeyde, Michael Elad & Matan Protter. *On single image scale-up using sparse-representations*. In Curves and Surfaces, pages 711–730. Springer, 2012.

[Zhang 02a] D. S. Zhang & G. J. Lu. *Generic Fourier descriptor for shape-based image retrieval*. IEEE International Conference on Multimedia and Expo, Vol I and Ii, Proceedings, pages 425–428 1549, 2002.

[Zhang 02b] D. S. Zhang & G. J. Lu. *Shape-based image retrieval using generic Fourier descriptor*. Signal Processing-Image Communication, vol. 17, no. 10, pages 825–848, Nov 2002.

[Zhang 03] Dengsheng Zhang & Guojun Lu. *Evaluation of MPEG-7 shape descriptors against other shape descriptors*. Multimedia Systems, vol. 9, no. 1, pages 15–30, 2003.

[Zhang 04] D. S. Zhang & G. J. Lu. *Review of shape representation and description techniques*. Pattern Recognition, vol. 37, no. 1, pages 1–19, Jan 2004.

[Zhang 05] Dengsheng Zhang & Guojun Lu. *Study and evaluation of different Fourier methods for image retrieval*. Image and Vision Computing, vol. 23, no. 1, pages 33–49, 2005.

[Zhang 08] Li Zhang, Yuan Li & Ramakant Nevatia. *Global data association for multi-object tracking using network flows*. In Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on, pages 1–8. IEEE, 2008.

[Zhang 11] Xiaoqun Zhang, Martin Burger & Stanley Osher. *A unified primal-dual algorithm framework based on Bregman iteration*. Journal of Scientific Computing, vol. 46, no. 1, pages 20–46, 2011.

[Zhu 96] S. C. Zhu & A. Yuille. *Region competition: Unifying snakes, region growing, and Bayes/MDL for multiband image segmentation*. IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 18, no. 9, pages 884–900, Sep 1996.

[Zibulevsky 10] M. Zibulevsky & M. Elad. *L1-L2 Optimization in Signal and Image Processing*. IEEE Signal Processing Magazine, vol. 27, no. 3, pages 76–88, May 2010.

[Zitova 03] B. Zitova & J. Flusser. *Image registration methods: a survey*. Image and Vision Computing, vol. 21, no. 11, pages 977–1000, Oct 1 2003.